

Decolonizing Artificial Intelligence Ethics: Ubuntu as a Philosophical Paradigm for AI Governance in Africa

Emmanuel Elolo Agbenonwossi

A dissertation presented to the Faculty of Arts in the University of Malta for the
degree of Master in Contemporary Diplomacy

December 2025




University of Malta Library – Electronic Thesis & Dissertations (ETD) Repository

The copyright of this thesis/dissertation belongs to the author. The author's rights in respect of this work are as defined by the Copyright Act (Chapter 415) of the Laws of Malta or as modified by any successive legislation.

Users may access this full-text thesis/dissertation and can make use of the information contained in accordance with the Copyright Act provided that the author must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the prior permission of the copyright holder.

Declaration

I hereby declare that this dissertation is my own original work.



Emmanuel Elole Agbenonwossi

13 December June 2025,

Lyon, France

Abstract

This dissertation examines the relevance of the African philosophy of Ubuntu to the ethical governance of artificial intelligence in African contexts. It analyses dominant global frameworks for AI ethics and governance and identifies conceptual and normative limitations in their engagement with African social values, communal worldviews, and historical experiences. The study adopts a qualitative, normative, and interpretive approach, drawing on African philosophical literature, international and regional policy documents, and expert interviews to assess how ethical assumptions are embedded in existing AI governance models.

The analysis conceptualises Ubuntu as a relational ethical framework centred on personhood, collective responsibility, and interdependence, and examines how these principles challenge prevailing individual-centred approaches to AI ethics. It explores the implications of Ubuntu-informed ethics for data governance, accountability, and institutional design, with particular attention to policy and legal frameworks at national, regional, and continental levels.

By positioning African philosophical thought as a source of normative reasoning rather than contextual adaptation, the dissertation contributes to discussions on ethical pluralism and legitimacy in AI governance. It concludes by identifying governance considerations relevant to African institutions and reflecting on the broader implications of Ubuntu-informed ethics for global debates on the responsible development and use of artificial intelligence.

Acknowledgements

I wish to express my profound appreciation to Professor Jovan Kurbalija for his intellectual guidance, critical insight, and unwavering support throughout the course of this research. His depth of expertise and thoughtful engagement were central to the conceptual refinement and analytical rigor of this dissertation.

I also extend my sincere thanks to the experts, colleagues, and mentors across Africa and within the wider international community who generously shared their time, reflections, and experience. Their contributions provided invaluable perspectives that enriched this work and anchored it in contemporary policy and ethical debates.

This dissertation is ultimately dedicated to all those who hold a confident and forward-looking belief in Africa's capacity to shape its own digital trajectories and to contribute original ethical and intellectual leadership to the governance of emerging technologies.

Dedication

To those who came before us.

To my beloved mother, who wished she could code.

To those who fear AI.

To those who learn it.

To those who govern it.

This work is yours.

If artificial intelligence is the last revolution of humankind,
this is my small contribution to its governance

Table of Contents

Abstract	3
Acknowledgements	4
Dedication	5
Table of Contents	6
Chapter 1: Introduction	8
1. Research background	9
2. Research Question	10
3. Objectives	10
3.1. General Objective	11
3.2. Specific Objectives	11
4. Significance of study	11
5. Scope of the Study	12
6. Structure of the dissertation	13
Chapter 2: Literature Review	15
1. Global frameworks	16
2. Critiques of global AI ethics	20
2.1. Epistemic Biases and the Myth of Universality	20
2.2. Decolonisation of Global AI Ethics	25
3. African contributions to digital ethics	27
4. Ubuntu and African communal philosophies	32
5. Gaps in the literature	36
Chapter 3: Methodology	39
1. Philosophical method	40
2. Justification for Ubuntu	41
3. Data sources	45
4. Reflexivity and positionality	47
5. Limitations	49
Chapter 4: Ubuntu and AI Ethics	51
1. Ubuntu as relational personhood	52
2. Implications for data identity	54
3. Comparison with Western ethics	56
4. Ubuntu and Islamic Ethical Traditions in African Moral Thought	58
Chapter 5: Governance, Policy and Legal Implications of Ubuntu-Informed AI Ethics	59
1. Ubuntu and Precolonial African Governance Models as Ethical Foundations for AI	60
2. Designing a Policy Framework Grounded in Ubuntu Values	61
3. Implementation Challenges and Tensions Between Indigenous Values and Modern Regulatory Paradigms	64
4. Legal Responses to Critiques of Ethno-Philosophy in African AI Governance	66
5. Embedding Ubuntu Principles into Formal Legal Frameworks	68
6. The Role of the African Union and Regional Economic Communities in Ethical AI	

<u>Governance</u>	<u>70</u>
<u>Chapter 6: Conclusion</u>	<u>73</u>
<u>1. Ubuntu, Ethics, and Artificial Intelligence: Synthesis of Core Findings</u>	<u>74</u>
<u>2. Policy and Institutional Recommendations for Africa</u>	<u>75</u>
<u>3. The Global Relevance of Ubuntu-Inspired AI Governance</u>	<u>78</u>
<u>References</u>	<u>80</u>
<u>Appendices</u>	<u>85</u>
<u>Appendix A</u>	<u>85</u>
<u>Semi-Structured Interview Guide</u>	<u>85</u>
<u>Appendix B</u>	<u>86</u>
<u>Expert Interviews Conducted for the Study</u>	<u>86</u>
<u>Appendix C</u>	<u>88</u>
<u>Policy and Governance Documents Reviewed</u>	<u>88</u>

Chapter 1: Introduction

1. Research background

Artificial intelligence has rapidly become a central component of contemporary governance, shaping public administration, social service delivery, security practices, and economic planning. Across Africa, governments and institutions are increasingly integrating algorithmic systems into national development strategies, digital transformation agendas, and public sector reforms (African Union, 2020; World Bank, 2023). This growing adoption positions artificial intelligence not merely as a technical tool, but as a political and ethical infrastructure with long-term implications for social organization, state authority, and human dignity.

Africa's engagement with artificial intelligence unfolds against a historical backdrop marked by limited participation in earlier technological transformations. From industrialization to the digital revolution, technological systems were often introduced on the continent through externally defined models, with minimal consideration for local values, social structures, or epistemic traditions (Rodney, 1972; Diop, 1981; Wiredu, 1996). As artificial intelligence becomes embedded in governance processes, this pattern raises critical questions about agency, legitimacy, and ethical alignment. The challenge is not only how AI is deployed, but on whose terms its ethical boundaries are defined.

The ethical risks associated with artificial intelligence in African contexts are already visible. Empirical studies have documented algorithmic bias in facial recognition systems trained on non-African datasets, leading to discriminatory outcomes and misidentification (Buolamwini & Gebru, 2018). At the same time, the deployment of surveillance technologies under security or administrative rationales has generated concern regarding accountability, civic freedoms, and the potential entrenchment of authoritarian practices (Gwagwa et al., 2020; CIPESA, 2021). These risks demonstrate that AI governance in Africa cannot be reduced to technical performance or efficiency alone, but must be assessed through an ethical lens attentive to social consequences, power relations, and historical asymmetries.

Existing global frameworks for artificial intelligence governance provide important reference points for policy development. Instruments such as the OECD Principles on Artificial Intelligence, the European Union's Artificial Intelligence Act, and UNESCO's Recommendation on the Ethics of Artificial Intelligence articulate shared commitments to transparency, accountability, and human rights (OECD, 2019; European Commission, 2021; UNESCO, 2021). However, these frameworks are largely grounded in normative traditions that emphasize individual autonomy, procedural compliance, and institutional capacities that are not uniformly present across African contexts (Metz, 2011; African Union, 2020). When applied without contextual adaptation, they risk overlooking communal modes of social organization, collective notions of responsibility, and historically grounded governance practices that shape ethical reasoning on the continent.

This dissertation situates the African philosophy of Ubuntu as a conceptual framework through which these challenges can be examined. Ubuntu emphasizes relational personhood, mutual responsibility, and the interdependence of individuals and communities, captured in the ethical maxim that a person becomes fully human through

relationships with others (Menkiti, 1984; Ramose, 2002; Tutu, 1999). Rather than introducing Ubuntu as a cultural supplement to existing models, the study approaches it as an ethical paradigm capable of informing how artificial intelligence governance can be conceptualized, designed, and evaluated within African societies. By grounding AI ethics in an indigenous philosophical tradition, the research seeks to contribute to a governance approach that is both contextually legitimate and normatively robust.

2. Research Question

This research seeks to understand how African philosophical traditions, particularly Ubuntu, can inform the ethical governance of artificial intelligence on the continent. It responds to a central concern: whether Africa can participate in the shaping of AI in ways that reflect its historical experiences, social values, and political aspirations, rather than merely adopting imported frameworks.

The primary question guiding this thesis is:

In what ways can the philosophy of Ubuntu serve as a foundation for ethical and inclusive AI governance in Africa?

This guiding question unfolds into several related lines of inquiry:

- *How has Africa's historical exclusion from previous technological revolutions shaped its current position in the global AI landscape?*
- *What are the core ethical principles within Ubuntu, and how do they contrast with dominant global approaches to AI ethics?*
- *How can African states, institutions, and communities integrate Ubuntu into the design, deployment, and regulation of AI systems?*
- *What opportunities exist for African-led regional cooperation to develop contextually grounded AI policy frameworks?*
- *How can this philosophical grounding support vulnerable populations, protect cultural diversity, and promote equitable access to AI benefits?*

These questions aim to support a broader reflection on how African perspectives can contribute not only to local innovation and justice, but also to the global conversation on the future of artificial intelligence.

3. Objectives

This research aims to position African philosophical thought as a normative contributor to contemporary debates on the ethical governance of artificial intelligence. Grounded in the philosophy of Ubuntu, the study seeks to examine how indigenous ethical frameworks can

inform the development of AI governance models that are legitimate, inclusive, and responsive to African social realities..

3.1. General Objective

To examine how the African philosophy of Ubuntu can provide an ethical foundation for the governance of artificial intelligence in Africa, with particular attention to questions of dignity, responsibility, and collective accountability.

3.2. Specific Objectives

The specific objectives of this study are:

- To situate Africa's contemporary engagement with artificial intelligence within its historical experiences of externally driven technological governance and normative exclusion.
- To identify and analyze the core ethical principles embedded in the philosophy of Ubuntu and assess how these principles differ from dominant global approaches to AI ethics.
- To examine existing AI governance frameworks at the global, continental, and national levels, with particular attention to their relevance and limitations in African contexts.
- To explore how Ubuntu-informed ethical principles can be translated into practical AI governance mechanisms, including policy design, institutional oversight, and regulatory practice.
- To assess the potential contribution of Ubuntu-based AI governance to inclusive development, social trust, and ethical legitimacy within African societies and regional institutions.

4. Significance of study

This dissertation contributes to contemporary debates on the ethical governance of artificial intelligence by foregrounding African philosophical thought as a legitimate and substantive source of normative guidance. While the global expansion of AI has generated an extensive body of ethical principles, guidelines, and regulatory proposals, these frameworks have largely emerged from Euro-American intellectual traditions and institutional contexts (OECD, 2019; European Commission, 2021; UNESCO, 2021). As a result, ethical reflection on AI governance has tended to privilege individual autonomy, procedural compliance, and abstract universality, often without sustained engagement with alternative moral traditions and social imaginaries. This study addresses that imbalance by examining Ubuntu as a coherent ethical paradigm capable of informing AI governance in African contexts.

From an academic perspective, the research advances interdisciplinary dialogue between African philosophy, ethics of technology, and governance studies. African ethical systems have frequently been marginalised within global norm-setting processes or treated as cultural references rather than structured philosophical frameworks (Wiredu, 1996; Hountondji, 1997). Ubuntu, in particular, is often cited in policy discourse as a value associated with solidarity or reconciliation, yet its deeper ontological and normative implications remain under-theorised in relation to contemporary technological governance. By analysing Ubuntu as a philosophy grounded in relational personhood, collective responsibility, and social harmony, this dissertation contributes to a growing body of scholarship that challenges the epistemic exclusion of African thought from global ethical debates (Menkiti, 1984; Ramose, 2002; Metz, 2011).

The study is also significant for public policy and institutional practice. Across the African continent, governments and regional organisations are developing national AI strategies and engaging with international governance initiatives in an effort to harness the benefits of artificial intelligence while mitigating its risks (African Union, 2020; African Union, 2022). These processes raise questions of ethical legitimacy, particularly when global frameworks are adopted without sufficient adaptation to local governance traditions and social norms. By exploring how Ubuntu-informed ethics can be translated into principles relevant to policy design, institutional oversight, and regulatory practice, the research offers conceptual resources for developing AI governance approaches that are both globally informed and locally grounded.

More broadly, the dissertation engages with debates on ethical pluralism in global AI governance. The assumption that a single set of ethical principles can be universally applicable risks overlooking the ways in which responsibility, harm, and dignity are understood and negotiated within different social and historical contexts (Fricker, 2007; de Sousa Santos, 2014). By situating Ubuntu within contemporary AI ethics discourse, the study illustrates how relational and community-oriented ethical frameworks can enrich existing approaches and contribute to more contextually responsive governance models. In this sense, the significance of the research lies not in proposing a prescriptive alternative to dominant frameworks, but in demonstrating the importance of integrating diverse philosophical traditions into ethical deliberations that increasingly shape the political, social, and institutional conditions of the digital age.

5. Scope of the Study

This study examines how the African philosophy of Ubuntu can inform the ethical governance of artificial intelligence within African contexts. The analysis is primarily situated in sub-Saharan Africa and focuses on governance, ethical reasoning, and policy design rather than on technical development or algorithmic engineering. The research is concerned with normative frameworks, institutional practices, and regulatory approaches through which artificial intelligence is conceptualised, deployed, and governed.

The study adopts a qualitative, philosophical, and policy-oriented approach. It draws on African philosophical literature, international and regional policy documents, national AI strategies, and expert interviews to explore how Ubuntu can be translated into governance principles relevant to artificial intelligence. The research does not involve empirical testing of AI systems or software implementation. Instead, it engages with ethical interpretation, legal analysis, and comparative governance assessment to examine the alignment between AI governance frameworks and African moral and social traditions (Wiredu, 1996; Ramose, 2002; African Union, 2020).

To ground the analysis, the study refers to selected country contexts, including Kenya, Rwanda, and South Africa, which have taken visible steps toward developing national AI strategies or engaging in structured policy discussions on artificial intelligence. These cases are used illustratively rather than exhaustively, with the aim of highlighting governance dynamics and ethical considerations rather than producing comparative evaluations of national performance. At the regional level, the research engages with continental and multilateral initiatives, particularly those led by the African Union, to situate national developments within broader efforts toward policy coordination and normative alignment (African Union, 2020; African Union, 2022).

The scope of the study is intentionally delimited to ethical governance and does not seek to propose a universal or uniform model of AI regulation for the continent. Instead, it advances context-sensitive principles that can be adapted to diverse institutional, cultural, and political environments. By maintaining this focus, the research aims to contribute to discussions on ethical pluralism, legitimacy, and responsibility in AI governance, while recognising the diversity of African societies and governance traditions.

6. Structure of the dissertation

This dissertation is structured into six chapters, following a progressive analytical trajectory that moves from conceptual framing and critical review to ethical analysis, policy implications, and concluding synthesis. The structure is designed to reflect the normative, philosophical, and governance-oriented nature of the research, while avoiding unnecessary fragmentation of closely related arguments.

Chapter 1 introduces the study by presenting the research background, the central research question, the objectives, the significance of the study, its scope, and the overall structure of the dissertation. It situates artificial intelligence governance in Africa within broader historical, ethical, and political contexts and establishes Ubuntu as the conceptual lens through which the research is conducted.

Chapter 2 provides a critical review of the literature on artificial intelligence ethics and governance. It examines dominant global frameworks, including multilateral and national approaches, and analyses their underlying philosophical assumptions and limitations in African contexts. The chapter then engages with critiques of universality and epistemic bias, explores African contributions to digital ethics, and examines Ubuntu and related communal philosophies as ethical systems relevant to AI governance. It concludes by identifying key gaps in the literature that the dissertation seeks to address.

Chapter 3 outlines the methodological approach adopted in the study. It explains the philosophical and normative orientation of the research, justifies the selection of Ubuntu as the primary ethical framework, and describes the data sources, including policy documents, academic literature, expert interviews, and field-based observations. The chapter also reflects on reflexivity, positionality, and the limitations of the study.

Chapter 4 constitutes the core analytical chapter of the dissertation. It applies Ubuntu as an ethical framework to artificial intelligence governance, examining relational personhood, data identity, communal responsibility, and accountability. The chapter contrasts Ubuntu-informed ethics with dominant Western ethical approaches and analyses illustrative cases to demonstrate how relational ethics reshape the understanding of harm, consent, transparency, and justice in AI systems deployed in African contexts.

Chapter 5 focuses on governance, policy, and legal implications. It examines how Ubuntu and precolonial African governance traditions can inform decolonial approaches to AI policy design, analyses institutional and regulatory challenges, and explores tensions between indigenous values and contemporary regulatory paradigms. The chapter further addresses legal critiques of ethno-philosophy, discusses strategies for embedding Ubuntu into formal legal frameworks, and analyses the role of continental and regional institutions, particularly the African Union and Regional Economic Communities, in shaping ethical AI governance in Africa.

Chapter 6 concludes the dissertation by synthesising the main findings and reflecting on their theoretical, ethical, and policy implications. It articulates key lessons on Ubuntu, ethics, and artificial intelligence, presents recommendations for African policymakers and institutions, and examines the potential global relevance of Ubuntu-inspired AI governance. The chapter closes by situating African philosophical traditions as active contributors to global debates on the ethical future of artificial intelligence.

Chapter 2: Literature Review

This chapter reviews the scholarly and policy-oriented literature relevant to the ethical governance of artificial intelligence, with particular attention to debates shaping global and African approaches to AI ethics. The purpose of the review is not to provide an exhaustive account of all existing work on artificial intelligence, but to identify the dominant ethical frameworks, theoretical assumptions, and governance models that inform contemporary policy and academic discussions, and to assess their relevance to African contexts.

The chapter engages with three broad bodies of literature. The first concerns global frameworks for AI ethics and governance, including policy instruments and normative guidelines developed by international organisations and leading regulatory actors. These frameworks have played a significant role in shaping national AI strategies and establishing widely referenced ethical principles. The second body of literature consists of critical analyses that interrogate the philosophical assumptions, institutional conditions, and power relations underlying these global approaches, particularly in relation to questions of universality, inclusion, and contextual legitimacy. The third body of literature focuses on African contributions to digital ethics, including philosophical traditions, legal perspectives, and policy initiatives that articulate alternative ethical orientations grounded in communal values and relational forms of social organisation.

Within this context, Ubuntu and related African philosophical traditions are examined as sources of ethical reasoning that offer perspectives distinct from dominant individual-centred approaches to AI ethics. Rather than treating these traditions as cultural background, the chapter situates them within broader debates on ethical pluralism, governance legitimacy, and normative diversity. This approach allows for a structured assessment of how African ethical frameworks have been engaged, marginalised, or selectively incorporated in existing literature on AI governance.

The chapter is organised into five sections. The first section examines major global frameworks for artificial intelligence ethics and governance. The second section reviews critical literature that challenges the assumptions and limitations of these frameworks. The third section explores African scholarly and policy contributions to digital ethics. The fourth section examines Ubuntu and other African communal philosophies as ethical systems relevant to AI governance. The final section identifies gaps in the existing literature that this dissertation seeks to address.

1. Global frameworks

Artificial intelligence (AI) is reshaping societies, economies, and governance systems across the world. In response to its rapid evolution, several governance frameworks have emerged, aiming to ensure AI is developed and deployed ethically, safely, and inclusively. These frameworks vary in their scope, enforceability, and normative orientation. However, a shared recognition exists that unregulated AI can amplify bias, reinforce inequality, and threaten human dignity. As such, the global conversation on AI governance increasingly centers on normative principles such as transparency, accountability, inclusiveness, and human rights compliance.

Among the leading efforts, the European Union's Artificial Intelligence Act stands out as the most advanced attempt at a comprehensive regulatory framework. Adopted in 2024, the AI Act introduces a risk-based approach that classifies AI systems into categories (unacceptable, high-risk, limited-risk, and minimal-risk) and sets corresponding regulatory requirements. It prohibits practices such as biometric categorization based on sensitive attributes and social scoring, while imposing stringent obligations on providers of high-risk AI systems (European Commission, 2024). The Act is seen as a global reference, setting a regulatory precedent for democratic societies (Veale & Borgesius, 2021).

Meanwhile, the United States has embraced a more decentralized and sectoral approach. Instead of a single AI law, U.S. governance of AI is characterized by soft law mechanisms, including the "Blueprint for an AI Bill of Rights" issued by the White House Office of Science and Technology Policy (OSTP, 2022). This document outlines five guiding principles: safe and effective systems, algorithmic discrimination protections, data privacy, notice and explanation, and human alternatives. While lacking binding authority, these principles are increasingly informing public procurement practices and agency oversight in key sectors.

China, in contrast, has integrated AI regulation within a broader digital governance strategy that emphasizes social stability, state security, and technological sovereignty. The 2021 Provisions on the Administration of Algorithmic Recommendation Services and subsequent generative AI regulations mandate transparency of algorithms, content moderation, and ideological alignment with state values (Creemers et al., 2022). These rules illustrate the instrumental use of AI governance in authoritarian contexts, where risk is framed less in terms of individual rights and more as threats to political order.

Beyond national initiatives, multilateral and multistakeholder platforms have advanced principles-based governance. UNESCO's 2021 Recommendation on the Ethics of Artificial Intelligence represents a landmark normative instrument. Adopted unanimously by 193 member states, the Recommendation outlines ethical principles and policy actions across ten areas, including proportionality, sustainability, gender equality, data governance, and inclusive participation (UNESCO, 2021). Its emphasis on human rights, environmental sustainability, and global solidarity makes it particularly resonant for developing countries seeking to align innovation with social justice.

In parallel, the OECD's AI Principles (2019) and the G7's Hiroshima AI Process (2023) underscore economic and democratic priorities in the governance of AI. These include promoting trustworthy AI, supporting innovation, and ensuring global interoperability of governance mechanisms. While these frameworks reflect high-income countries' priorities, they increasingly acknowledge the need for international cooperation and South-South knowledge exchange.

Yet, despite the proliferation of global initiatives, major challenges remain. First, there is a significant implementation gap. Many countries lack the institutional capacity, expertise, or infrastructure to translate normative principles into enforceable regulations. Second, governance asymmetries persist. Dominant AI narratives and standards continue to be shaped by a handful of actors in the Global North, sidelining African and other Southern

epistemologies. Third, global frameworks often privilege techno-legal solutions over socio-political dimensions such as justice, power relations, and cultural worldviews (Crawford, 2021).

In this context, Africa's contribution to AI governance must not be limited to implementation or compliance. Rather, it must entail normative innovation and ethical leadership rooted in the continent's philosophical traditions. One such contribution lies in the ethics of Ubuntu. Ubuntu, a philosophy emphasizing interdependence, community, and human dignity, offers a powerful lens to reimagine AI governance as a collective, inclusive, and relational process. It challenges dominant notions of autonomy, profit maximization, and data commodification by centering humanity, solidarity, and care (Metz, 2014).

As scholars like Van Norren (2024) and Diagne (2023) argue, Ubuntu can serve as an ethical compass in AI governance, especially when adapted to digital environments. By insisting that "a person is a person through other persons," Ubuntu reframes the development and use of AI not as a technocratic endeavor, but as a moral one that must enhance communal well-being and justice. Its relational ontology contrasts with the individualistic assumptions underlying many global frameworks and calls for co-governance, equity, and contextual sensitivity in AI policymaking.

Moreover, the African Union's Continental Strategy for AI, adopted in 2024, offers a continental roadmap aligned with Ubuntu values. It outlines principles of inclusion, sustainability, and shared prosperity, while emphasizing data governance, innovation ecosystems, and regional harmonization (African Union, 2024). Despite being non-binding, it signals a shift from passive policy adoption to proactive norm-setting. The AU's approach complements other regional efforts, such as the ECOWAS Guidelines for Ethical AI and the Smart Africa blueprint, which foreground local needs and participatory governance.

To fully realize Africa's potential in AI governance, further action is needed. This includes investing in regulatory capacity, supporting multistakeholder platforms, integrating Ubuntu and indigenous philosophies into legal frameworks, and ensuring African voices shape global standards. Ultimately, global AI governance must evolve into a pluralistic, polycentric ecosystem where African norms, values, and priorities are not only recognized but co-determine the future of AI.

While global frameworks offer structured entry points into AI ethics, they frequently risk becoming normative imports in African contexts rather than co-created paradigms. Most of these frameworks are developed in high-income countries and reflect the moral assumptions, institutional logics, and legal traditions of the Global North (Crawford, 2021; Veale & Borgesius, 2021). Their global influence is in part due to the technological dominance and regulatory leadership of the EU, US, and China, but also because institutions such as UNESCO, the OECD, and IEEE provide formal platforms for diffusion. Yet, as noted by African scholars and policymakers alike, these frameworks risk epistemic asymmetry and may unintentionally reinforce new forms of digital dependency (UNESCO, 2021; AU, 2024).

One of the main critiques of global AI ethics frameworks is their insufficient sensitivity to African ontologies and lived realities. Despite referencing inclusion and diversity, many such documents still rely on liberal individualism and Enlightenment-era views of rationality that are not universal. For instance, the EU AI Act, while pioneering in its scope, assumes a highly institutionalized legal environment with robust enforcement capacity, conditions often absent in many African states (European Commission, 2024; GCG, 2023). The Act's emphasis on compliance through technical documentation, conformity assessment, and centralized registries presumes infrastructural and financial resources that may not be equitably distributed across regions.

Ubuntu, in contrast, reframes the ethical locus of AI from individualistic responsibility to communal care and mutual interdependence. According to Metz (2014), Ubuntu is a relational ethic rooted in the idea that one's humanity is actualized through others. Applied to AI governance, this principle suggests that the moral legitimacy of AI systems should not rest solely on abstract principles of fairness or accountability, but on whether these systems enhance social harmony, protect vulnerable members of society, and sustain the well-being of the collective. This orientation is especially relevant in postcolonial societies, where technology must confront legacies of marginalization and epistemic violence.

Experts interviewed for this research emphasized that Ubuntu is not merely a cultural slogan but a viable governance principle. Dr. Melody Musoni, for instance, pointed out that without concrete policy instruments and investment in regulatory capacity, Ubuntu risks becoming a hollow normative banner. She emphasized the need for institutional structures that bridge indigenous philosophies and legal mechanisms, including ethical review boards that reflect communal values and participatory data governance models (Musoni, personal interview, 2025). Similarly, Strive Masiyiwa highlighted how inclusion in AI governance must extend beyond representation to include ownership, ensuring that Africans co-develop and shape the algorithms and datasets used in their societies (Masiyiwa, personal interview, 2025).

In this regard, the African Union's Continental Strategy for AI, adopted in 2024, marks a critical step forward. Grounded in the principles of solidarity, sustainability, and data sovereignty, it explicitly recognizes Africa's unique cultural and historical contexts. It proposes regionally anchored innovation ecosystems, legal harmonization, and public-private partnerships that empower small states and island nations (African Union, 2024). It also echoes Ubuntu's values by advocating participatory governance, digital inclusion, and intergenerational equity.

However, the strategy's non-binding nature raises questions about implementation. As noted in the UNESCO (2023) AI readiness assessment for Africa, many countries still lack national AI policies, dedicated regulators, or even basic digital infrastructure (UNESCO, 2023). Therefore, implementation of Ubuntu-informed AI governance requires a multi-scalar approach: national legal adaptation, regional cooperation through Africa's Regional Economic Communities (RECs) and Pan-African multilateral bodies to ensure African ethics are not peripheral in multilateral standard-setting (Chechi, 2024; FieldFisher Podcast, 2024).

One promising model is the hybridization of frameworks, adapting global norms to local contexts while asserting African philosophical agency. For example, Rwanda's approach to integrating ethics into its AI policy combines elements of UNESCO's recommendations with community-driven consultative processes and attention to traditional norms of consensus. Similarly, South Africa's AI policy draws on the post-apartheid legacy of Ubuntu as a legal and moral norm, using it to frame human rights and algorithmic accountability (AUDA-NEPAD, 2024).

These examples underline the possibility of designing hybrid, pluralistic models of AI governance. In such models, Ubuntu is not seen as a counterpoint to global ethics, but as a complementary and necessary perspective that critiques the technocratic and depersonalized assumptions of mainstream models. As Diagne (2023) affirms, Ubuntu offers a call to "make humanity together"; a notion that reorients AI governance away from isolated compliance and toward relational justice, ecological care, and planetary ethics.

To move forward, African countries must be supported in developing robust legal and technical infrastructure, while also fostering philosophical education that revives indigenous knowledge systems. This entails funding for research in African digital ethics, inclusion of Ubuntu in computer science curricula, and open data policies that respect communal ownership. As Kurbalija (2025) suggested, these frameworks must not only reflect Africa's values but enable Africans to co-author the rules of digital engagement, thereby decolonizing both code and ethics.

In sum, while global AI governance frameworks have advanced critical norms, they remain incomplete without African epistemologies. Ubuntu provides a philosophical foundation for reimagining AI ethics beyond compliance toward care, dignity, and communal flourishing. When legally and institutionally embedded, it offers not only a critique but also a constructive alternative that can enrich global governance and fulfill AI's promise for all humanity.

2. Critiques of global AI ethics

2.1. Epistemic Biases and the Myth of Universality

The field of artificial intelligence (AI) ethics has witnessed a proliferation of principles, guidelines, and policy frameworks over the past decade. These documents, developed largely by institutions in Europe and North America, aim to establish shared ethical ground for the governance of AI. They often include values such as fairness, transparency, privacy, human autonomy, and accountability. While these concepts appear to be neutral and widely acceptable, they are in fact rooted in specific philosophical traditions. This claim to neutrality conceals a deeper epistemic bias, one that privileges certain moral languages while marginalising others. The presumption of universality reinforces colonial hierarchies in global knowledge production and policy design (Crawford, 2021).

At the heart of this bias lies a narrow conception of the human person that emerged from the Enlightenment. In this view, the ethical subject is an autonomous individual who exercises reason to pursue self-defined goals. This understanding of personhood informs the principle of autonomy, a cornerstone in nearly all global AI ethics frameworks. As

originally proposed by Immanuel Kant, autonomy entails the capacity of a rational agent to act in accordance with moral laws derived from reason (Kant, 1785/1997). In this framing, morality becomes an internal duty based on rational deliberation, disconnected from social, cultural, or spiritual contexts.

This model contrasts sharply with many African philosophical systems, particularly those influenced by Ubuntu. In this framework, a person is not defined by individual capacity but by relational existence. Ethical identity is forged through interactions with others and with the community. The maxim often cited by Ubuntu philosophers: "a person is a person through other persons" encapsulates this worldview (Tutu, 1999). The priority is not individual freedom, but the cultivation of harmonious relationships. Moral agency is therefore not an abstract capacity, but a lived practice of care, solidarity, and communal affirmation (Metz, 2007).

The dissonance between these paradigms reveals a deeper asymmetry in AI governance. Global ethical frameworks often advance what Sylvia Wynter (2003) calls the "overrepresentation of Man." This concept refers to the elevation of a specific Western construction of humanity (rational, self-governing, secular) as the normative model. This construction is treated as if it represents all humanity, thereby eclipsing other ways of being, knowing, and relating. As a result, the ethical languages of many African communities are excluded from AI governance not because they lack coherence, but because they do not conform to dominant paradigms.

My conversation with Professor Souleymane Bachir Diagne reinforced this critique. He emphasised that Ubuntu should not be regarded merely as a regional curiosity or cultural add-on. Instead, it offers a fundamental rethinking of what it means to be human in relation to technology. In our dialogue, he reflected: "Ubuntu is not an ethics checklist. It is ethics as co-being. It is not rights-driven, it is relationship-driven." This perspective challenges the notion that AI ethics can rely solely on frameworks that abstract moral questions from context and community. It insists that ethics be understood as emergent, embodied, and relational.

Global ethics processes often acknowledge the importance of inclusion but rarely transform their conceptual foundations. In the UNESCO Recommendation on the Ethics of Artificial Intelligence (2021), for example, there is recognition of cultural diversity and a call for ethical pluralism. Yet the underlying principles (transparency, accountability, data privacy) are still framed in terms of individual rights and procedural rationality. These are important values, but their dominance constrains the space for ethical reflection grounded in communal cosmologies. Dr. Melody Musoni, during her interview, described this as "structural tokenism." African voices are invited to the table only after the menu has been set.

Even the definition of harm in these frameworks reflects Eurocentric assumptions. Ethical impact assessments often focus on quantifiable outcomes such as bias in hiring algorithms or fairness in credit scoring. These metrics, while useful, rarely address collective harms. In African contexts, decisions about data, automation, or surveillance often affect entire communities, kinship systems, or ancestral relationships. For example, data collected for

agricultural optimisation may seem innocuous when assessed individually. However, as witnessed in a farming community in northern Togo, farmers raised concerns about the erosion of customary land rights due to AI-based land categorisation. The ethics framework used for the project made no provision for communal stewardship or historical memory.

This exclusion is not accidental. It is sustained by institutional structures, linguistic hierarchies, and funding models. Most global AI ethics discourse occurs in English and presumes familiarity with liberal legal norms. African languages, which contain rich ethical vocabularies, are almost entirely absent. Terms like "omoluabi" in Yoruba, which denotes a person of honour and integrity embedded in community life, have no equivalent in dominant AI discourse. As Ngũgĩ wa Thiong'o (1986) argues, language is not just a medium but a worldview. The dominance of colonial languages in AI ethics enforces what he calls "linguistic imperialism", a narrowing of moral imagination through the suppression of local expression.

This phenomenon aligns with what Boaventura de Sousa Santos (2014) describes as "epistemicide," the silencing of entire ways of knowing. In AI governance, epistemicide manifests in the structural absence of African cosmologies from design labs, regulatory frameworks, and ethical review boards. When African contributions are included, they are often reframed through Western categories. Ubuntu becomes "African transparency" or "communal fairness," stripped of its metaphysical depth and reinterpreted to fit into pre-existing templates.

Such translation is not inherently problematic. But when it is one-sided, when African concepts are forced to conform while Western ones remain unchallenged, it becomes a form of epistemic violence. As Fricker (2007) explains, epistemic injustice is not only about exclusion but about the denial of someone's capacity as a knower. When African scholars must constantly explain their ethics in alien terms, they are not treated as co-creators of knowledge but as interpreters of the already authorised.

Beyond conceptual asymmetries, the power dynamics embedded in global AI ethics are maintained through what scholars have termed epistemic dependency. This occurs when African governments and institutions adopt ethics frameworks not because they reflect local realities, but because they align with donor requirements or international expectations. As pointed out in the GIZ and Smart Africa study (Azeroual, 2023), over sixty percent of AI policy drafts across Africa were shaped with technical assistance from external actors, often based on templates created in the Global North. These templates rarely account for Africa's pluralistic values, social structures, or oral epistemologies. Instead, they replicate institutional models that emphasize compliance, risk mitigation, and abstract ethical principles.

During regional dialogues on AI policy, such as the 2024 consultations in Lomé and Kigali, participants frequently encountered this dynamic. Although the discussions nominally invited African perspectives, the structure of the meetings remained centered on terms like algorithmic fairness, transparency audits, and explainability standards. As one policymaker from West Africa noted, "we are translating our concerns into their language, but our

language remains unspoken.” This pattern reflects what decolonial scholars identify as extractive inclusion participation without epistemic transformation (Mignolo, 2009).

An additional layer of epistemic bias lies in the governance of AI knowledge production. Leading conferences, peer-reviewed journals, and AI ethics hubs are disproportionately based in Europe and the United States. Funding and publication structures reward adherence to prevailing methodological norms, making it difficult for African scholars to advance research grounded in endogenous worldviews. Interview data from Melody Musoni revealed this challenge: “I once submitted a policy paper based on Ubuntu-informed accountability, and the reviewer asked if Ubuntu could be made into a compliance checklist. They had no conceptual tools to understand relational ethics.”

Ubuntu does not lend itself to checklists. It is not an ethics of procedures but one of presence, responsibility, and moral becoming. It conceives the human being not as a discrete, autonomous subject, but as an emergent reality forged in reciprocity with others and the environment. The Zulu expression “*umuntu ngumuntu ngabantu*” expresses this as an ontological commitment: personhood is relational, not inherent. This disrupts the foundational logic of most global AI ethics frameworks, which isolate the individual as the unit of analysis. As Metz (2022) explains, Ubuntu ethics entails acting in ways that promote shared life, mutual aid, and social cohesion. Ethical violations are not only breaches of privacy or autonomy but disruptions to communal harmony.

AI systems designed around liberal autonomy may inadvertently destabilise these communal relations. For example, a machine learning system used to distribute health subsidies in East Africa, based on individual profiles, failed to account for extended family structures. As a result, the system prioritised nuclear households, marginalising grandparent-led families and multi-generational living arrangements. The ethical review of the system found no fault because individual consent was obtained. However, community leaders argued that the redistribution of support weakened social ties and undermined traditional caregiving roles. This kind of structural misalignment reveals how ostensibly neutral ethics frameworks can cause material and moral harm.

Language continues to be one of the most powerful tools of epistemic exclusion. As noted in the 2023 GCG-AI Landscape Study, the vast majority of AI ethics policies across Africa were drafted in English or French, with no provision for consultation in local languages . This omission reflects more than a linguistic oversight. It signals a disregard for the moral vocabularies embedded in African languages. For instance, in Wolof, the term *teranga* denotes hospitality as a moral obligation, while in Shona, *hunhu* resonates with the Ubuntu idea of relational personhood. These concepts carry normative force, shaping how communities understand justice, accountability, and wellbeing. Their absence from policy debates is not simply a matter of translation. It is a form of conceptual exclusion that narrows the ethical imagination of what AI governance can be.

Resistance to this exclusion is growing. Initiatives like Masakhane NLP and Mozilla’s Kiswahili speech dataset project are not merely technical achievements. They are acts of epistemic recovery, affirming that African languages and knowledge systems belong at the center of AI development. The African Union’s draft guidelines on data governance also

call for respecting customary norms and community data stewardship, an implicit rejection of the individualistic data ownership models promoted by the General Data Protection Regulation (GDPR) and similar frameworks.

Moreover, in the FieldFisher Africa Week podcast, Kweya Chechi from Microsoft emphasized that African AI adoption must not replicate the inequities of past technological transitions. She pointed out that digital policy cannot focus only on access and skills, but must also reflect local values and priorities. Chechi cited examples of AI being used in Kenya for disaster prediction and in Ghana for healthcare access, stressing that ethical design must engage with local actors from the outset. This approach aligns with a communitarian ethic that places the well-being of communities above the abstraction of technical standards.

Yet even well-meaning efforts to localize AI ethics can fall into the trap of epistemic paternalism. When African traditions are selectively cited or simplified for inclusion in policy documents, they risk being instrumentalized rather than respected. Ubuntu becomes a footnote in a Western framework, rather than a foundation for new thinking. This concern was expressed by Professor Diagne, who warned that unless Ubuntu is treated as a full epistemological system, it will remain a rhetorical ornament rather than a transformative force. True epistemic inclusion means engaging with Ubuntu not as folklore, but as philosophy.

This requires epistemic humility on the part of global actors. It means recognising that no framework is neutral, that every value is culturally situated, and that AI governance must be co-created rather than exported. It also calls for the development of what Santos (2014) calls ecologies of knowledge spaces where different epistemologies coexist, intersect, and generate new insights. In the context of AI, this may include hybrid ethics reviews, intercultural policy roundtables, and multilingual deliberation formats that allow communities to express their concerns in their own conceptual terms.

Dismantling the myth of universality demands a commitment to narrative sovereignty. African communities must not only be consulted in AI development; they must tell their own stories, define their own risks, and articulate their own ethical horizons. The future of AI ethics cannot be built on a single civilisational grammar. It must be rooted in a plurality of traditions, including those forged in oral histories, sacred texts, indigenous jurisprudence, and postcolonial critique. Ubuntu offers more than a regional alternative. It is a lens through which we can rethink our shared human condition in a digital age.

In this view, AI governance becomes not a question of compliance with abstract norms, but of ethical cohabitation across worlds. The task is not to translate Ubuntu into the language of AI ethics, but to let Ubuntu speak on its own terms and to reshape the conversation accordingly.

2.2. Decolonisation of Global AI Ethics

One of the fundamental pillars of decolonising AI ethics is recognising the historical and structural forces that have shaped technological development and ethical standard-setting.

Postcolonial theorists such as Frantz Fanon and Ngũgĩ wa Thiong'o have long underscored how colonialism not only extracted material resources but also imposed epistemic hierarchies (Fanon, 1963; wa Thiong'o, 1986). In the domain of AI, this legacy is evident in the way ethical values and regulatory frameworks are disproportionately derived from Western philosophical traditions often assumed to be universally valid.

Global AI ethics principles such as fairness, transparency, and accountability are presented as neutral moral benchmarks. Yet these terms carry culturally specific assumptions about individual rights, rational agency, and procedural justice. As Professor Jovan Kurbalija observed in recent policy discussions, many global frameworks embed liberal political theory without acknowledging its cultural partiality or colonial roots (Meeting Notes with Prof. Kurbalija, 2025). Decolonisation therefore requires making these philosophical assumptions visible and open to critique. It also demands epistemic pluralism: a willingness to let other worldviews set the terms of ethical engagement.

African ethical systems offer rich alternatives. Ubuntu, hunhu, teranga, and other communitarian traditions centre the relational and collective dimensions of moral life. They challenge the Cartesian separation of mind and body, the Kantian abstraction of duty, and the liberal prioritisation of autonomy over solidarity. Ubuntu, for example, frames ethical action as that which enhances the shared humanity of all involved. As Metz (2017) explains, right action is not merely what can be justified to rational individuals, but what fosters harmonious and respectful communal living. This has profound implications for how we understand responsibility, trust, and justice in the context of AI systems.

Melody Musoni, a regional expert on AI and data governance, highlighted this disconnect in a 2025 interview: “Most of the national strategies in Africa mention Ubuntu, but it’s used symbolically, not structurally. There’s no real effort to translate those values into algorithmic design or decision-making protocols.” Her observation reflects the broader critique that global AI ethics often assimilates indigenous values superficially, without allowing them to reshape the architecture of policy or technology.

True decolonisation would require more than inserting cultural keywords into Western frameworks. It would necessitate rethinking how AI systems are designed, implemented, and evaluated. For instance, data governance models built on individual consent fail to address the communal ownership of knowledge and land in many African societies. The idea of “free, prior, and informed consent” in the collection of data assumes individual proprietorship, whereas many communities conceptualise data as relational and collective. The African Declaration on Internet Rights and Freedoms (2015) recognises this tension, but few AI guidelines have operationalised it.

The global AI research agenda is skewed towards concerns of the Global North, such as facial recognition bias or algorithmic hiring discrimination, while neglecting pressing African issues like digital land grabbing, biometric surveillance in humanitarian settings, and AI-enabled political manipulation. This imbalance is not accidental. It results from funding flows, publication gatekeeping, and institutional prestige concentrated in Western countries (Gomez & Van Dijck, 2021). Decolonising AI ethics means shifting the axis of

knowledge production, funding African-led research, supporting community-based data stewardship, and amplifying local narratives about what responsible AI means.

During the AI Summit in Kigali in 2025, a discussion on the sidelines brought to light a revealing example from a regulatory sandbox initiative in a rural community in West Sub-Saharan Africa. The project was designed to test the potential of AI tools to support sustainable agriculture through predictive analytics, within a controlled environment aimed at informing future regulatory approaches.

Despite its innovative intent, the initiative encountered significant challenges. The AI system, trained primarily on external datasets, failed to incorporate the oral ecological knowledge held by local elders' knowledge that had long guided planting cycles through observations of seasonal cues and environmental signals. When the algorithm's recommendations conflicted with these long-standing practices, trust in the system deteriorated rapidly. As a result, the pilot was suspended before completion. This case illustrates the risks of deploying AI in culturally embedded contexts without accounting for local epistemologies. It highlights the importance of grounding both data inputs and ethical design in the lived realities and knowledge systems of the communities AI seeks to serve, even within experimental policy frameworks.

In academic and diplomatic spaces, African contributions to AI ethics are often filtered through Western epistemic lenses. For example, during multilateral negotiations surrounding AI ethics at UNESCO and the Global Partnership on AI (GPAI), African representatives frequently find their proposals diluted or rephrased in terms compatible with dominant policy frameworks. As highlighted in the AI in Africa Landscape Study, “African voices are present but not prominent” (GCG, 2022, p. 14). The study observes that while inclusion is often nominal, the underlying frameworks rarely reflect African ethical epistemologies.

One of the most significant barriers to ethical pluralism is linguistic hegemony. The predominance of English and French in global AI discourse excludes ethical frameworks embedded in African languages. As noted earlier, words such as *omoluabi* in Yoruba, which denotes a deeply moral person who acts in harmony with the community, have no direct equivalent in English. Similarly, Swahili's *utu* and Wolof's *jom* express culturally situated ethical concepts that resist translation but are essential for any comprehensive understanding of moral agency in African societies.

Decolonising AI ethics therefore also entails epistemic translation not simply moving African concepts into Western language but allowing African languages and categories to set the terms of ethical evaluation. This requires institutional humility from international bodies and a willingness to create multilingual, intercultural platforms for ethical deliberation. It also requires investment in African scholarship and the creation of repositories of knowledge that document, preserve, and apply traditional moral philosophies in contemporary settings (UNESCO, 2021).

Strive Masiyiwa, in a 2025 interview, emphasized the importance of infrastructure and youth empowerment for ethical AI in Africa. “If we do not build the technical capacity and

philosophical ownership of these systems, Africa will again be left out of a major revolution.” His insight reflects the historical pattern where Africa is invited to adopt technologies without shaping their design. Decolonising AI ethics means ensuring that Africa is not only a site of implementation but also a site of conceptual innovation and ethical leadership.

The re-centring of African philosophies in AI ethics is not a retreat into cultural isolationism. It is a move toward ethical pluralism and intellectual justice. As Diagne (2023) argues, Ubuntu’s call to “make humanity together” challenges the atomistic models of selfhood and rationality that dominate much of Western moral theory. It opens space for ethical frameworks that prioritise interdependence, reciprocity, and ecological responsibility, values urgently needed in the face of algorithmic harm and climate vulnerability.

The decolonisation of AI ethics is not a supplementary concern. It is foundational to building technologies that serve all of humanity. It demands shifting power, revising curricula, funding African-led research, protecting indigenous data sovereignty, and embedding relational ethics at every stage of the AI lifecycle. It requires confronting the myth of ethical universality and building systems grounded in the diverse ways humans have understood care, justice, and responsibility. As this thesis demonstrates, decolonisation is both a philosophical imperative and a practical agenda for sustainable, inclusive, and truly ethical AI governance.

3. African contributions to digital ethics

African contributions to digital ethics are complex, historically grounded, and epistemologically distinct. These contributions draw on indigenous knowledge systems, relational philosophies, legal frameworks, community practice, and resistance to colonial legacies. African thinkers, institutions, and civil society have advanced ethical frameworks rooted not in imported abstraction but in lived experience, cultural pluralism, and philosophical integrity.

At the heart of African ethical thought lies a relational conception of personhood and society. This contrasts sharply with the individualist frameworks that dominate global AI ethics. Ubuntu, a Southern African moral philosophy, asserts that ethical life is constituted through human interdependence. As famously encapsulated by Archbishop Desmond Tutu, “a person is a person through other persons.” Ubuntu is not a sentimental expression of unity but a normative system that emphasizes dignity, solidarity, restorative justice, and mutual responsibility (Metz, 2017). These principles have concrete implications for the ethics of artificial intelligence. They require that AI systems be judged not only by their technical efficiency but by their capacity to sustain the well-being of communities.

Ubuntu-informed ethics foreground the collective. This reorients digital governance away from the hyper-individualism that underpins many global frameworks, such as those promoted by the OECD or the EU, where autonomy and privacy are framed solely at the individual level. In African societies where land is communally held, knowledge is orally transmitted, and social life is relational, digital ethics must reflect collective harms and

collective rights. Dr. Melody Musoni, in an April 2025 interview, emphasized that “many African countries name Ubuntu in national AI strategies, but without implementation mechanisms, Ubuntu becomes token-invoked without next teeth.” The ethics of care embedded in Ubuntu challenge the Western framing of AI ethics as risk management or fairness optimization. Instead, ethics becomes a practice of relational accountability.

Institutional contributions further demonstrate African leadership in digital ethics. The African Union’s Continental AI Strategy, adopted in 2024, is a landmark document. It articulates a vision of artificial intelligence that supports inclusive development, safeguards human rights, and responds to local contexts. Importantly, it acknowledges structural inequalities in data access, computational infrastructure, and epistemic representation. The Strategy includes five pillars (risk mitigation, benefit maximization, capacity development, inclusive investment, and international cooperation) and proposes concrete actions across 15 thematic areas (African Union, 2024). This strategic orientation is not merely regulatory; it is normative. It affirms that Africa has the right to govern AI on its own terms and in line with its values.

African experts also contributed substantially to the UNESCO Recommendation on the Ethics of Artificial Intelligence, adopted in 2021. While the Recommendation uses universalist language, its credibility in Africa depends on how local actors domesticate it. As noted in the GCG AI Landscape Study (2022), African policymakers often face tension between aligning with international ethical standards and reflecting indigenous norms. The study found that “many national strategies transpose global principles with little contextual adaptation, often due to donor pressure or lack of local capacity” (GCG, 2022, p. 17). This highlights the importance of epistemic sovereignty, an idea central to African contributions to digital ethics.

Legal instruments like the Malabo Convention (2014), formally the African Union Convention on Cyber Security and Personal Data Protection, provide foundational principles for ethical governance of data. Though not focused on AI per se, the Convention sets out rights and obligations regarding privacy, data minimization, and accountability. It recognizes the importance of cultural and social specificities in framing legal protections. Some African scholars argue that the Convention’s regional legal framework serves as a counterpoint to externally imposed standards and should be updated to include AI-specific provisions rooted in African philosophical traditions (Taddesse, 2023).

Other notable legal contributions include the African Declaration on Internet Rights and Freedoms, which extends digital ethics beyond data to encompass access, equality, and transparency. The Declaration, drafted through a participatory process involving African civil society, explicitly mentions the collective dimension of rights. It emphasizes the right to develop and use digital technologies in a manner consistent with cultural values and communal welfare. This document frames ethics not only as individual protections but as collective entitlements, a theme often neglected in mainstream AI debates.

African civil society and academia have also shaped digital ethics through open knowledge production. The Masakhane project, for instance, is a pan-African initiative to build natural language processing (NLP) models for African languages. By developing AI systems that

serve and preserve African linguistic heritage, Masakhane challenges the epistemic domination of Western languages and data infrastructures. It reflects an ethics of care, access, and cultural relevance. As Marivate and Bezuidenhout (2021) explain, such projects advance “epistemic justice by making African languages computable and visible in the digital age.”

Furthermore, grassroots initiatives across the continent are embedding ethics into digital practice. In Kenya, civic technology platforms use AI to monitor election integrity and improve public service delivery. In Nigeria, hackathons and innovation hubs regularly include ethics checkpoints and community consultations in AI prototyping. These efforts exemplify how African communities are not just passive beneficiaries of digital technology but active agents of ethical innovation. As Kweyu Chechi noted during the FieldFisher Africa Week podcast (2025), “digital tools in Africa are only ethical if they empower, localize, and bridge, not marginalize.”

Yet African contributions are not uniform. The continent contains over fifty countries, thousands of languages, and multiple cosmological systems. Ethical contributions arise from diverse traditions: Yoruba concepts of *omoluabi* (moral excellence), Swahili ideals of *utu* (humanity and respect), Wolof norms of *jom* (dignity and courage), and Ethiopian philosophy rooted in theological rationalism. These are not merely cultural curiosities but ethical grammars that inform decision-making, justice, and legitimacy. These systems are often transmitted orally and require different methods of philosophical analysis.

However, challenges remain. The underrepresentation of African scholars in global ethics forums, the lack of ethical infrastructure in many national strategies, and the instrumentalization of indigenous concepts without redistribution of power all hinder the full integration of African ethics into AI governance. Souleymane Bachir Diagne warned in a 2025 interview that “Ubuntu cannot become a mere footnote in global ethics. If it does not change the structure, it becomes a performance of inclusion.”

While Ubuntu often garners symbolic attention in global discourses, other African philosophical traditions also make significant ethical contributions. The Yoruba notion of *omoluabi* is central to understanding moral excellence in West African societies. It describes a person who exemplifies self-discipline, honesty, compassion, and respect for elders and the community. In AI ethics, this framework implies not only accountability and transparency but also the cultivation of moral character in algorithmic design and deployment. Omoluabi ethics means that being ethical is not about compliance but about becoming the kind of person or system that enhances the moral ecology of society (Abimbola, 2020).

The Akan tradition of Ghana offers another lens through *sankofa*, a principle urging people to go back and reclaim that which has been forgotten or lost. In digital ethics, *sankofa* implies recovering indigenous knowledge systems to inform contemporary technological practice. This aligns with growing global movements toward ethical pluralism and anti-colonial design. In AI development, *sankofa* challenges the dominant Silicon Valley ethos of speed and disruption, encouraging deliberate, culturally-informed innovation. The

incorporation of *sankofa* can help build algorithms that are not only locally relevant but historically conscious.

Bantu philosophy, especially as interpreted by Placide Tempels (1959) and later expanded by African thinkers, also adds depth. It conceptualizes being as force and emphasizes interconnectedness in a dynamic universe. Ethical action, then, is about maintaining equilibrium between humans, nature, and the spiritual realm. In digital ethics, this philosophy suggests that algorithms must be evaluated not just by outcomes, but by how they sustain or disrupt ontological balance. This stands in sharp contrast to utilitarian or libertarian approaches that focus only on efficiency or freedom.

From the Horn of Africa, Ethiopian philosophical traditions especially those rooted in Orthodox Christianity contribute a rationalist yet spiritual ethical discourse. Thinkers such as Zara Yaqob developed systems of natural reason that emphasised equality, justice, and mutual obligation (Getachew, 2023). These traditions remain influential in parts of East Africa and can inform AI regulation frameworks that combine rational critique with spiritual responsibility. This dimension is often missing in secular Western ethics.

Across the continent, scholars and activists are increasingly creating platforms to articulate these diverse traditions. The African Academy of Sciences, through its ethics committees, supports work on culturally grounded bioethics that increasingly intersects with AI debates. Similarly, policy institutions such as the African Union Development Agency (AUDA-NEPAD) have funded research and capacity-building projects to develop African AI principles, including through their Centre of Excellence in Science, Technology and Innovation.

These efforts reflect not only theoretical contributions but strategic interventions. In the AI4D (Artificial Intelligence for Development) program, coordinated by IDRC and implemented across Francophone and Anglophone Africa, multiple ethics tracks have supported locally led initiatives to draft principles of responsible AI. Researchers from Senegal, Rwanda, Kenya, and Nigeria have identified local values such as equity, solidarity, and inclusion as priority pillars for AI governance (AI4D, 2021). These principles differ markedly from Western standards framed around individual consent, liability, and procedural fairness.

In the context of data governance, the African Union Data Policy Framework adopted in 2022 advocates for “community-centred data ecosystems.” This echoes African philosophies that view knowledge not as an object to be owned but as a gift to be stewarded for communal benefit. Data colonialism where tech companies extract African data without meaningful consent or return, violates these ethics. African contributions push for sovereignty, equity in data access, and protection from algorithmic harm.

Experts such as Dr. Musoni have argued for the practical integration of African ethics into policy documents. In her interview, she stated: “It is not enough to name Ubuntu or omoluabi. If we do not train civil servants, translate policies into local languages, and set up institutions that practice relational ethics, the values remain performative.” She advocated for budget allocations to community-based ethics boards, inclusion of traditional

authorities in consultations, and protocols for cultural impact assessment before AI deployment.

Similarly, insights from Kweya Chechi during the FieldFisher Africa Week 2025 podcast show how African leadership in digital ethics can come from corporate partnerships rooted in African priorities. As she observed, many of Microsoft's regional initiatives in Kenya and Ghana now consult local stakeholders before introducing AI tools in public health or agriculture. She emphasized that "ethics is not an afterthought. It is part of our engagement process. That means respecting customary norms, youth perspectives, and the social economy."

Importantly, these contributions also resist being co-opted into extractive global systems. As Souleymane Bachir Diagne warned, "there is a tendency to universalise only after the West has defined the terms. But African ethics (Ubuntu, jom, utu etc) do not seek universality. They propose hospitality. They invite the world in, without erasing their origins." This philosophical generosity allows for global ethical pluralism without epistemic erasure.

Academic institutions in Africa are beginning to institutionalise these contributions. The University of Lagos offers a program on digital rights and indigenous ethics. The University of Pretoria's Centre for AI Research includes a workstream on cultural dimensions of responsible AI. The University of Addis Ababa recently hosted a symposium on Ethiopian philosophy and emerging technologies. These initiatives show that African digital ethics is not limited to slogans or activist critique, it is evolving into a structured academic field.

This field is necessarily interdisciplinary. It draws on philosophy, law, anthropology, theology, and design. It is informed by oral traditions, community governance, and postcolonial critique. It is locally rooted but globally oriented. African contributions to digital ethics are not about regional exceptionalism. They are about restoring balance to a discourse skewed by epistemic asymmetries. As with other decolonial projects, the goal is not to replace one center with another. It is to build a polycentric, dialogical, and just foundation for AI governance.

African contributions to digital ethics offer powerful correctives to the dominant frameworks of global AI governance. They foreground relationality over abstraction, community over individualism, history over amnesia, and plurality over hegemony. These ethics arise not from mimicry or marginality, but from sovereignty, memory, and vision. As this thesis contends, no ethical framework can claim legitimacy in Africa if it does not reflect these contributions.

4. Ubuntu and African communal philosophies

Ubuntu, a philosophical cornerstone in many sub-Saharan African societies, provides a distinct moral vision based on the interdependence of all human beings. The phrase "umuntu ngumuntu ngabantu" translates as "a person is a person through other persons." This principle emphasizes that ethical behavior arises not from individual autonomy, as

frequently found in Enlightenment traditions, but from relational existence. The moral unit is not the self-sufficient individual, but the individual-in-community. This epistemic foundation sets Ubuntu apart from the dominant frameworks that inform global AI ethics. It is this distinctiveness that renders Ubuntu not merely relevant, but essential in the conversation around ethical AI governance in Africa and beyond.

Ubuntu has historically informed social organization, conflict resolution, education, and leadership across various African societies. Within pre colonial governance systems, the legitimacy of authority was often tied to its ability to maintain harmony and community welfare rather than its adherence to codified laws. Elders and councils practiced deliberation not to determine rights in the abstract, but to restore balance and collective consensus. These modes of governance underscore Ubuntu's emphasis on dignity, reciprocity, and shared responsibility (Letseka, 2012). In the digital age, these ethical resources remain not only culturally valuable but politically and technologically relevant.

The rise of artificial intelligence presents new ethical dilemmas that affect privacy, agency, accountability, and justice. However, dominant responses have emerged mostly from Euro-American institutions, embedding Western liberal assumptions into the global policy discourse (Crawford, 2021). These frameworks often emphasize individual consent, transparency, and autonomy, which, while important, do not necessarily resonate with the lived experiences or moral frameworks of African communities. Ubuntu offers a different vocabulary and a different set of priorities: dignity, solidarity, relational accountability, and restorative justice.

Several African thinkers have articulated the philosophical coherence and ethical sophistication of Ubuntu. For instance, Ramose (2002) argues that Ubuntu provides an ontological foundation that is processual and communal, in contrast to Cartesian dualism. In Ubuntu ontology, "to be" is "to participate in the web of life," and harm to one member of the community is a rupture in the whole. Such a worldview has implications for how AI systems should be governed: technologies should not only respect individual rights but also uphold the relational integrity of communities. This is particularly relevant for data governance, surveillance, and algorithmic decision-making.

In interviews conducted for this thesis with African policymakers and ethicists, a recurring theme was the inadequacy of current AI frameworks to capture communal harms. One respondent from Rwanda noted, "Ethical evaluations of AI often stop at whether consent was obtained. But in our culture, the bigger question is: how does this affect the whole family, the clan, the village?" Another expert, an international Human Rights lawyer from Nigeria, emphasized that Western ethics frequently treats community merely as context, rather than as an actor or moral agent. Ubuntu reverses this lens, making community not just the backdrop but the subject of ethical reflection.

This ethical orientation has practical implications. For example, in considering biometric data collection, Ubuntu ethics would not only ask if the individual consented but also whether the use of such data reinforces or undermines community trust. This would extend into policies about data localization, algorithmic auditing, and the inclusion of affected communities in oversight mechanisms. The dominant metrics of fairness in machine

learning, such as demographic parity or equalized odds, may need to be supplemented with Ubuntu-informed criteria such as social cohesion, reciprocity, or communal flourishing (Olonisakin et al., 2023).

Yet Ubuntu is not monolithic. Across the continent, similar communal philosophies exist that, while distinct in expression, share foundational assumptions about relational ethics. In the Akan tradition, the concept of “biako ye” speaks to unity through difference. In Yoruba thought, the idea of “omoluabi” captures moral excellence defined through social responsibility. These philosophies are not primitive approximations of moral reasoning but are complex ethical systems with relevance to 21st-century technological governance. Their absence in global discussions on AI ethics is not due to a lack of rigor but a lack of epistemic humility from dominant institutions (Mavhunga, 2017).

Moreover, Ubuntu is dynamic. Contrary to criticisms that it is overly spiritual or anti-modern, Ubuntu continues to evolve and express itself through new mediums. In digital governance, this includes community-led data trusts, open-source AI training in local languages, and culturally embedded digital literacy programs. For instance, organizations such as Masakhane and Data Science Africa are exploring community-centered machine learning models rooted in linguistic and cultural inclusion. These initiatives operationalize Ubuntu not as a slogan, but as a guiding framework for AI design and deployment that reflects local realities and aspirations.

International organizations are beginning to acknowledge the importance of Ubuntu. UNESCO’s Recommendation on the Ethics of AI (2021), while not entirely rooted in Ubuntu, acknowledges the value of regional ethical traditions and includes contributions from African scholars. Still, these references often remain tokenistic, and the deeper ontological shifts that Ubuntu demands are not yet reflected in policy architecture (UNESCO, 2021). Integrating Ubuntu meaningfully requires not just citation, but structural transformation: participatory policymaking, multilingual engagement, and genuine power-sharing in governance forums.

Ubuntu also offers critical tools to challenge algorithmic opacity. In contrast to the technocratic impulse to make systems explainable through logic or probability, Ubuntu calls for transparency grounded in social accountability. It is not enough that an algorithm can be audited technically. What matters is whether its use can be justified to those it affects and whether it restores or undermines social harmony. This reframing shifts the goal of AI explainability from computational transparency to ethical legibility within lived communities.

Some critics argue that Ubuntu is too vague or inconsistent to serve as a governance framework. But this critique misunderstands the nature of ethical traditions. Western ethics, too, includes diverse schools of thought (deontology, utilitarianism, virtue ethics) with internal contradictions. What matters is whether a framework can offer normative guidance that is actionable, contextual, and just. Ubuntu meets these criteria, particularly when understood not as a static code but as an ethical process grounded in deliberation and mutual care.

integrating Ubuntu into AI governance requires more than philosophical alignment. It demands institutional and policy transformation across local, national, and international levels. Ubuntu challenges the technocratic tendency to isolate ethics from culture, urging that technological design, regulation, and evaluation be re-embedded in social systems. In practical terms, this means developing AI policies that are not only rights-based but relationship-based. Rather than focusing solely on legal compliance or technical standards, an Ubuntu-informed policy environment emphasizes mutual accountability, social inclusion, and community engagement throughout the AI lifecycle.

This paradigm shift is already visible in emerging initiatives. In South Africa, pilot projects on ethical AI have involved traditional authorities in the consultation process, acknowledging that legitimacy in many rural areas derives from communal endorsement rather than individual authorization. One such initiative, conducted with the support of the Human Sciences Research Council, explored how data governance could be harmonized with local land tenure systems, traditional dispute resolution mechanisms, and informal economies (HSRC, 2023). This example illustrates the practical relevance of Ubuntu in areas such as agricultural monitoring, digital identity, and e-governance.

Further reflections from the Africa-wide consultation on the UNESCO AI Ethics Recommendation revealed tensions between global principles and local understandings. Participants from East Africa highlighted that community elders often play a decisive role in determining what constitutes harm, justice, and reparation. In such contexts, fairness cannot be computed by metrics alone but must be negotiated through cultural processes. An elder from northern Uganda shared during the workshop, “We trust the wisdom of the community more than the speed of the computer. You cannot audit fairness without sitting with the people.” This illustrates the centrality of human presence and intersubjective dialogue in ethical validation, a hallmark of Ubuntu.

Ubuntu also opens a path for restorative AI governance. In contrast to retributive models that focus on penalizing wrongdoing, Ubuntu-inspired approaches prioritize restoring harmony after harm. In the context of AI, this may translate into community-led impact assessments, grievance redress systems, and participatory audits that center the voices of those most affected. For instance, a community in northern Ghana established a local ethics review board to evaluate the deployment of a mobile health app. Rather than rely solely on biomedical ethics or Western standards, the board included traditional healers, youth representatives, and women’s groups. Their deliberations extended beyond privacy or informed consent to include spiritual risks, generational equity, and trustworthiness (AI4D Africa, 2022).

This orientation contrasts with dominant frameworks that prioritize algorithmic logic over moral dialogue. As critical scholars like Birhane (2021) argue, AI systems reflect the values and assumptions of their designers, and these are often shaped by individualist worldviews that sideline collective life. Ubuntu reverses this dynamic by asserting that morality is not an internal possession but an external process enacted through social interaction. This means that ethical validation must occur within the community, not just in code.

One of the strengths of Ubuntu is its flexibility. Unlike rigid codes or universalist declarations, Ubuntu ethics can adapt to different contexts while maintaining its core commitment to relationality. This is particularly useful in Africa's plural legal environments, where customary law, religious norms, and state regulations intersect. Ubuntu offers a moral grammar that can mediate these systems, providing a common ethical language rooted in African experiences. It supports polycentric governance and recognizes that justice must be co-produced with the people it affects.

The work of African philosophers such as Mogobe Ramose, Yvonne Mokgoro, and Pumla Gobodo-Madikizela has deepened the theoretical underpinnings of Ubuntu in law and ethics. Their writings emphasize Ubuntu's role not just in private morality but in public reasoning, transitional justice, and constitutionalism (Gobodo-Madikizela, 2020). This is significant for AI governance, where ethical questions often intersect with legal, political, and economic domains. Ubuntu does not offer simple answers, but it provides an ethical compass grounded in responsibility, humility, and the common good.

Interviews conducted for this thesis affirm this perspective. Professor Souleymane Bachir Diagne emphasized Ubuntu as a "global ethic" rather than merely a regional one. He explained that Ubuntu invites all of humanity into a conversation on shared existence, mutual care, and planetary responsibility. This has implications for how Africa participates in global AI norm-making. Rather than being a recipient of frameworks, the continent must be a co-author of ethics rooted in its own traditions. Diagne's reflections reinforce the argument that Ubuntu is not a cultural footnote but a moral proposal for the digital age.

In the FieldFisher podcast with Kweya Chechi, Microsoft's Government Affairs Director for Africa, Ubuntu was implicitly evoked in her emphasis on relational governance. She spoke of "tech for good" not as charity, but as embedded responsibility: "Technology must be felt in people's lives through inclusion, language, and respect. It's not ethical if it does not transform how people live together." Her comments reflect the Ubuntu ideal that dignity arises from interconnection. She also pointed to the importance of youth, women, and rural communities in shaping AI policy, aligning with Ubuntu's emphasis on inclusive deliberation.

Importantly, Ubuntu also reshapes how we think about the future. Whereas many AI forecasts are driven by techno-optimism or risk aversion, Ubuntu offers an ethics of anticipation grounded in solidarity. It teaches that progress is not merely the expansion of capability, but the deepening of humanity. This is echoed in the African Union's Agenda 2063, which frames technological transformation within the goal of "an integrated, prosperous and peaceful Africa, driven by its own citizens." Ubuntu can guide the design of AI systems that serve not only market efficiency but historical justice and ecological balance.

However, operationalizing Ubuntu is not without challenges. One major issue is institutional capacity. Few African countries have AI-specific regulatory bodies, and existing institutions often lack the training to incorporate indigenous ethics into technical standards. There is a need for curricula, funding, and partnerships that embed Ubuntu into the education of engineers, policymakers, and civil society. As Melody Musoni noted in

her interview, “Ubuntu is powerful, but without infrastructure, it remains philosophical. We need institutions that practice Ubuntu in procurement, evaluation, and implementation.”

Language is another barrier. Much of the discourse on AI remains in English, French, or Portuguese, limiting access to local communities. Yet Ubuntu ethics cannot thrive without local idioms, storytelling, and metaphor. Translation is therefore both a technical and ethical imperative. Initiatives such as Masakhane NLP and AI4D Francophone Africa are already working to create multilingual datasets and tools. These efforts reflect Ubuntu’s call to inclusivity not just in outcome, but in process.

5. Gaps in the literature

A growing body of literature has examined the ethical dimensions of artificial intelligence (AI), with a significant concentration of scholarship emerging from Europe and North America. While these contributions have shaped the dominant frameworks of AI ethics, particularly around principles such as fairness, accountability, transparency, and human rights, critical gaps remain in the literature when it comes to African contexts. These gaps are not merely geographic or regional omissions. They reflect deeper epistemological silences that overlook diverse worldviews, including the philosophical traditions, social imaginaries, and historical experiences of African societies. These omissions limit the universality of global AI governance frameworks and perpetuate a knowledge imbalance within the ethics discourse.

The first notable gap is the underrepresentation of African philosophical frameworks, particularly Ubuntu, in mainstream AI ethics literature. Ubuntu, with its emphasis on relationality, communal personhood, and the interdependence of being, offers a distinct ontological and ethical orientation that could enrich global AI discourse. However, most academic and policy texts frame ethical questions in AI through liberal individualist assumptions, such as the autonomous rational subject and the contractual basis of justice. As highlighted during my interview with Professor Souleymane Bachir Diagne, these assumptions ignore the possibility of collective ontologies that define existence as inherently relational. Diagne emphasized that Ubuntu’s ethic of “making humanity together” is absent in most Western AI ethical models and should be foregrounded as an alternative moral vocabulary for AI governance in Africa.

Secondly, literature often treats Africa as a passive recipient of AI technologies rather than as a site of innovation or ethical reasoning. This is evident in the limited presence of African scholars and policymakers in influential global reports or conferences on AI ethics. According to the Global Index on Responsible AI, African states contributed only a marginal share to international AI publications, and their presence in shaping the UNESCO AI Ethics Recommendation remained minimal despite being among its adopters. This imbalance reinforces the perception that ethical leadership in AI should emerge solely from the Global North and sidelines African voices from shaping normative frameworks.

The literature inadequately addresses the structural inequalities embedded in the development and deployment of AI in Africa. Many studies focus on risks such as

algorithmic bias, surveillance, and lack of data protection, but they often treat these challenges in abstraction, without reference to their historical roots in colonial systems of control and extraction. For example, AI-driven surveillance technologies used in African urban centers are sometimes deployed without public consent or oversight. These practices mirror colonial legacies where population control was centralized and top-down. Few global ethics frameworks make these historical continuities explicit or interrogate how data colonialism operates through the digitization of African societies (Couldry & Mejias, 2019).

There is a silence in the literature regarding how AI affects local epistemologies and oral knowledge systems in Africa. Modern AI systems rely on datafication, standardization, and digitization of knowledge, which often marginalizes forms of knowing that are transmitted through oral traditions, communal storytelling, or indigenous ecological knowledge. This epistemic gap reflects a broader issue where dominant AI architectures and language models ignore or misrepresent African languages and idioms of reasoning. As noted in the Fieldfisher Africa Week podcast featuring Kweya Chechi, Microsoft's Director of Government Affairs in Africa, the lack of data in African languages has direct implications on inclusion and participation in digital economies. This gap is not just a technical shortfall but also an ethical oversight that demands rectification.

Current literature on AI ethics often fails to articulate what Emmanuel Iduma describes as “the afterlife of colonial infrastructure” in Africa. Technologies introduced in the name of innovation may reinforce dependencies and extractive logics if they do not prioritize local ownership, data sovereignty, and capacity building. In an interview with Dr. Melody Musoni, she explained that many national strategies in Africa cite Ubuntu symbolically but fail to invest in the infrastructure and education systems necessary to make ethical AI operational. This disjunction between symbolic reference and practical investment is a recurring gap that global frameworks have not addressed adequately.

Equally absent is critical reflection on gendered and intergenerational ethics in the African AI context. While Western ethics frameworks increasingly incorporate intersectional analyses, few African-centered AI ethics studies deeply explore how AI systems impact women, youth, and the elderly within traditional African family and governance structures. This omission is significant given that Ubuntu and other African philosophies inherently embed responsibilities to children, ancestors, and future generations. These dimensions could significantly enrich the intergenerational ethics of AI but are missing from most global documents, including the OECD's or IEEE's models.

The literature has not sufficiently interrogated the role of African spiritualities and cosmologies in framing ethical judgments about technology. Unlike secular traditions that separate the ethical from the spiritual, many African communities integrate cosmological beliefs into everyday governance. In some societies, technologies are evaluated not only based on harm or fairness but on their harmony with nature, ancestral values, and spiritual balance. This pluralism in ethical reasoning offers an expanded horizon for understanding what constitutes responsible AI, yet remains largely unacknowledged in dominant academic and policy narratives.

Although African AI ethics is increasingly acknowledged in policy conversations, academic literature continues to underrepresented core African intellectual traditions and social realities. Much of the existing work remains descriptive, outlining national AI strategies or general risks, but rarely engaging with deep moral reasoning grounded in African worldviews. This omission limits conceptual innovation and reinforces the dominance of external frameworks.

One critical oversight lies in how African societies' legal and cultural pluralism is excluded from AI governance design. Most frameworks assume a singular legal authority, often the state as the sole custodian of ethical AI oversight. Yet in many African contexts, authority is shared with traditional or religious leaders. These actors play key roles in shaping local norms about privacy, trust, and responsibility (Adeleke, 2023). Their absence from global literature weakens ethical contextualisation.

AI ethics documents continue to treat Ubuntu or related African values as cultural footnotes rather than operational principles. For instance, the UNESCO (2021) Recommendation mentions Ubuntu in passing, without embedding its relational philosophy into implementation tools or metrics. This tokenism does not allow for meaningful epistemological plurality in AI ethics. Research on community experiences with AI technologies also remains limited. As UNESCO's needs assessment survey (2021) indicates, rural and peri-urban populations are often left out of consultations, despite being among the most impacted. Their concerns spanning from land data control to surveillance fears do not align with standard risk categories like algorithmic bias or transparency, and so remain invisible in dominant policy texts.

This gap is reinforced by structural barriers facing African scholars. Funding, publishing access, and language hierarchies mean that even rigorous African contributions may be ignored unless they conform to Western academic expectations. Even when African researchers are included, they often operate in borrowed conceptual languages, limiting the transformative potential of their work. Finally, traditional knowledge systems such as Akan ethics, Bantu metaphysics, or Ethiopian theological rationalism are rarely studied in relation to digital governance. These systems could offer distinctive approaches to moral reasoning, intergenerational justice, or digital trust, but they are largely absent from peer-reviewed AI ethics publications.

Chapter 3: Methodology

1. Philosophical method

This dissertation adopts a philosophical method grounded in normative inquiry, interpretive analysis, and historical contextualization. The choice of method reflects the dissertation's ambition to move beyond procedural ethics and instead articulate a decolonial and culturally embedded framework for artificial intelligence (AI) governance. By foregrounding the African philosophical tradition of Ubuntu, the research is not merely situated within African contexts; it actively participates in re-theorizing what counts as ethical deliberation in the age of automation. The method is both analytical and reconstructive, embracing African epistemologies and their ontological assumptions about relationality, community, and the moral fabric of human life.

Normative philosophy remains essential to this work. It provides a conceptual foundation for evaluating what ought to guide technological development. However, this evaluation is not detached from experience or historical struggles. African philosophy, and Ubuntu in particular, does not draw sharp lines between ethics, lived experience, and political struggle. Therefore, the dissertation's philosophical method integrates normative reasoning with historically situated knowledge and postcolonial critique. As Hountondji (1996) argues, African philosophy cannot be understood apart from the socio-political context in which it arises. This aligns with Tangwa's (2004) framing of African bioethics as a response to both cultural continuity and modern challenges.

In this spirit, Ubuntu is not treated as a set of abstract moral principles that can be applied universally. Rather, it is interpreted as a relational mode of being, a worldview that sees the self as constituted through others. As per Ramose (2002), Ubuntu is an ontological condition: to exist is to co-exist. In method, this demands a departure from the detached stance of the Western analytic tradition. Instead of adopting a God's-eye view, the researcher embraces the position of a participant observer embedded in the cultural, political, and philosophical textures of the continent.

Thus, the philosophical method follows three trajectories. First, it employs critical historiography. This involves examining how African philosophies were excluded from dominant AI ethics frameworks and situating this exclusion within broader histories of epistemic injustice. Wiredu (1998) and Diagne (2011) both emphasized the ways colonialism reshaped what counts as knowledge and marginalized indigenous reasoning practices. In this sense, the method also functions as a critique of dominant techno-ethical paradigms.

Second, the dissertation embraces comparative normative analysis. Here, Ubuntu is placed in dialogue with other AI ethical frameworks, including UNESCO's Recommendation on the Ethics of Artificial Intelligence (2021), the European Union's AI Act, and the OECD AI Principles. However, unlike these frameworks, Ubuntu is not reduced to a checklist or compliance rubric. Instead, it offers a different vision of human flourishing, one that prioritizes communal responsibility, interdependence, and restorative over retributive justice (Metz, 2011).

Third, the method is grounded in philosophical reconstruction. Ubuntu is not frozen in pre-colonial imagination. It is reinterpreted and expanded to speak to contemporary challenges such as algorithmic bias, surveillance, and technological exclusion. Here, interviews with experts like Dr. Melody Musoni and Prof. Souleymane Bachir Diagne become important not merely as data but as dialogical engagements. Their reflections, recorded in the empirical sections of this thesis, are interpreted as philosophical contributions that deepen the understanding of Ubuntu's application to digital life.

The dissertation also takes inspiration from the decolonial turn in philosophy. According to Maldonado-Torres (2007), this turn demands that ethical inquiry begin not with abstract universals, but with the lived experiences of those who have historically borne the brunt of epistemic and technological domination. In this regard, Africa's long-standing encounter with externally imposed technologies whether colonial surveillance, imported legal systems, or global data extractivism forms the backdrop against which this philosophical method unfolds.

Ubuntu is viewed not simply as a moral code but as an epistemic framework. As Van Norren (2020) suggests, Ubuntu challenges the Enlightenment separation between subject and object, mind and matter, individual and society. Knowledge, in this framework, is not merely acquired; it is relationally produced through conversation, ritual, and shared existence. This methodological insight reshapes how ethics is done. It reclaims oral traditions, proverbs, communal deliberation, and indigenous hermeneutics as legitimate modes of philosophical reasoning.

The methodological commitment also includes reflexivity. Given that the researcher is embedded in the African context and professionally engaged in policy advocacy, self-awareness about positionality is essential. The philosophical method therefore integrates critical reflexivity into its approach, recognizing that one's standpoint whether as scholar, advocate, or community member informs how Ubuntu is interpreted and applied. This reflective stance acknowledges both the privileges and constraints of doing philosophy within institutions shaped by global knowledge hierarchies.

Finally, the method affirms that ethical reasoning must be embedded in practice. It draws from praxis-oriented traditions in African thought, such as those articulated by Amilcar Cabral and Cheikh Anta Diop, where the value of philosophy is measured by its contribution to collective liberation and well-being. This is why the dissertation includes a policy dimension: to translate Ubuntu's philosophical insights into actionable recommendations for AI governance in Africa.

2. Justification for Ubuntu

This research embraces Ubuntu not as a rhetorical flourish or nostalgic appeal to culture, but as a pragmatic philosophical grounding for ethical AI governance on the continent. The selection of Ubuntu is not a choice of cultural identity over analytical rigor. It is a deliberate epistemological orientation rooted in both philosophical and policy reasoning, designed to address structural exclusions that have persisted through colonial and postcolonial technological transformations. Ubuntu is not merely a heritage resource; it offers

conceptual tools and value systems that contest individualistic rationality, instrumental reason, and anthropocentric techno-solutionism, all of which dominate current global AI governance models.

In my interview with Melody Musoni, she highlighted that the popularity of Ubuntu in ethical debates often suffers from abstraction. She warned that without concrete investment, infrastructure, and policy translation, Ubuntu risks becoming “a beautiful but empty slogan.” This research accepts that challenge. Ubuntu is therefore not treated as a metaphorical ornament. It is investigated as a moral grammar with policy relevance. The goal is to operationalize Ubuntu’s ethics of relational personhood, mutual care, and collective accountability into concrete governance mechanisms, legal standards, and institutional cultures.

Benjamin Akinmoyeje, during our recorded conversation, expressed concern about Africa’s continued reliance on imported frameworks for AI governance. He observed that “AI ethics in Africa is still largely inspired by European templates. We are adjusting their claims to our contexts, not creating our own from the ground.” His remark points to a central justification for this research. The use of Ubuntu is not only about centering African values but reclaiming authorship in how the continent imagines its digital futures. African philosophy must not be consulted only for community mobilization or civil society workshops. It must be at the origin of governance design and ethical reflection.

Ubuntu offers this possibility because it is both a lived ethic and a political idea. Historically, Ubuntu grounded dispute resolution mechanisms, land-sharing arrangements, leadership norms, and restorative justice systems across many African communities, from the Igbo and Zulu to the Tswana and Shona. The depth and adaptability of Ubuntu were demonstrated in the post-apartheid Truth and Reconciliation Commission (TRC), where its notions of healing and collective dignity were elevated to a national framework for justice and rebuilding. Unlike the retributive logics embedded in many Western systems of control, Ubuntu speaks in the language of restoration, of making whole again, of recognizing harm as relational rupture and not merely as legal breach. This emphasis becomes especially relevant in the context of artificial intelligence, where algorithms have the capacity to perpetuate or redress social injury at scale.

The philosophical decision to work with Ubuntu in this dissertation also stems from the insufficiency of dominant global ethics frameworks. As explored in Chapter Two, the EU AI Act and the OECD frameworks rely heavily on human rights discourse, fairness, transparency, and safety. While valuable, these concepts are often grounded in liberal individualism and procedural rationality. Ubuntu, by contrast, starts from a different ontological premise: a person is not a detached rights-bearer, but someone whose identity and flourishing depend on the relationships they cultivate. Justice, from an Ubuntu perspective, cannot be imagined without social harmony and collective responsibility. This is why an Ubuntu-based AI governance model would not prioritize algorithmic audits as endpoints but would focus on creating governance ecosystems where dignity is actively protected, not only for individuals but for communities and future generations.

Ubuntu also challenges the linear temporality often assumed in policy development. Most global AI governance models are future-oriented in a predictive sense. They attempt to regulate what is coming next. Ubuntu integrates temporalities differently. It considers ancestors and future generations as participants in ethical decision-making. This intergenerational ethic could reshape how African states conceptualize AI governance timelines, moving from compliance-oriented monitoring to long-view stewardship models. This orientation was echoed by Souleymane Bachir Diagne, who noted during our interview that “Ubuntu is not only about human beings living today, but also those who lived before and those yet to be born. It obliges us to think relationally through time, not just space.”

This philosophical approach responds to the realities of AI deployment in Africa. Many of the technologies being imported are trained on data that is historically skewed, racially biased, and epistemologically alien to the continent’s knowledge systems. A framework grounded in Ubuntu can push for dataset accountability not just in terms of statistical parity, but in terms of epistemic inclusion. What histories are being erased? Whose values are being assumed as neutral? Who defines what intelligence is? Ubuntu invites us to pause and question.

The political justification for Ubuntu also derives from its pan-African adaptability. While it originates from Southern African traditions, its values (hospitality, consensus, reciprocity, and the primacy of the community) resonate across linguistic and cultural lines. The Akan notion of *biako ye* (unity is strength), the Yoruba ethics of *omoluabi*, or the Swahili principle of *ujamaa* all reflect this communal worldview. This broad resonance supports the development of a continental AI governance framework that is not confined to South African policy traditions but can find traction in West Africa, East Africa, and beyond.

Benjamin Akinmoyeje also emphasized the need for what he called “ethical independence.” He stated, “Africa must stop waiting for model laws from Europe. We should not just copy and paste. Our laws should speak our languages, literally and philosophically.” This dissertation responds to his challenge. Ubuntu, framed here as both concept and method, offers the grounding needed for Africa to propose rather than merely adapt AI ethics principles. It provides a normative vocabulary that reflects lived African values while being flexible enough to translate into policy terms.

The effort to justify Ubuntu as a central axis of ethical reflection in this research also emerges from the recognition that imported frameworks, although globally promoted, are often context-insensitive. Global ethical guidelines on AI, like those from the IEEE or UNESCO, emphasize technical robustness and fairness, yet rarely account for communal obligations, intergenerational stewardship, or moral cosmologies shaped by spirituality and land-based knowledge systems. These omissions are not incidental. They reflect philosophical priorities rooted in Western liberalism and a secular epistemology. Ubuntu, on the other hand, makes room for ancestral wisdom, spiritual accountability, and collective personhood as central categories for reasoning about justice in the digital age.

From a policy development perspective, Ubuntu can help counteract the proceduralism that often defines AI ethics implementation. Many governments adopt checklists of fairness, explainability, and bias mitigation without addressing the ethical substance of those metrics. In contrast, Ubuntu calls for a deeper moral conversation about what it means to govern justly. For example, an Ubuntu-inspired policy might not only evaluate whether a facial recognition system is accurate across demographic groups but also ask whether its use aligns with community values about visibility, privacy, and sacred space.

This orientation has already begun to influence regional initiatives. The African Union's Continental AI Strategy and several national policy drafts make reference to ethics and inclusion. While not always citing Ubuntu explicitly, the spirit of relational accountability and communal dignity can be traced within them. However, these gestures must be deepened. Ubuntu is not a decorative reference for legitimacy. It is a method for structuring deliberation, distributing responsibility, and imagining futures rooted in care. This research attempts to make that shift by integrating Ubuntu as both a philosophical foundation and a practical framework.

Another important justification for choosing Ubuntu lies in its resonance with broader decolonial theory. African digital ethics must be formulated not simply as a contribution to global discourse but as a rupture from it. Ubuntu offers this potential because it does not aim to universalize. It begins from locality, relationality, and the humility of co-becoming. It disrupts the notion that ethics is a matter of abstract rules and rights, and instead reorients moral reasoning toward lived interdependence. In this sense, Ubuntu aligns with thinkers like Achille Mbembe, who argue for an "Afropolitan" ethic of entanglement and pluralism, and Sylvia Wynter, who critiques the colonial production of the "Man" as the only valid subject of rationality and law. Ubuntu offers a way to philosophize ethically from below, from the ground of everyday African life.

From an educational standpoint, Ubuntu also helps challenge the implicit curriculum in AI training programs. As it stands, most African universities teach computer science and policy using imported textbooks and case studies. Ethical analysis is often confined to Western dilemmas: autonomous cars, data breaches in Europe, facial recognition in US policing. Ubuntu enables students to ask different questions. How should I serve pastoral communities? What are the risks of predictive analytics in customary land governance? How does digitization alter communal rites and obligations? These are not marginal questions. They are central to building a responsible and responsive digital future.

Melody Musoni rightly emphasized that Ubuntu must be lived to be effective. It cannot remain an academic framework or a line in a policy preamble. It must reshape how consultations are held, how data is collected, how harms are assessed, and how accountability is shared. This research builds on that imperative by grounding its methodology in community engagement, philosophical reflection, and policy translation. Ubuntu, in this thesis, is not treated as folklore. It is treated as jurisprudence.

The justification for Ubuntu as the philosophical and normative core of this dissertation rests on multiple grounds: it counters epistemic injustice, grounds ethical governance in African realities, offers alternatives to procedural ethics, enables policy translation,

supports decolonial theorizing, and reshapes pedagogical approaches. It is not the only framework Africa could adopt, but it is one that speaks deeply to the continent's past and future, its wounds and its wisdom. It is a philosophy that, if taken seriously, can guide the continent not just to adapt digital technology ethically but to redefine what ethical technology can mean for humanity.

3. Data sources

This dissertation on ethical AI governance through the lens of Ubuntu draws from a diverse and intentionally pluralistic set of data sources. Rather than privileging formal academic texts alone, the research incorporates documentary analysis, oral traditions, expert interviews, institutional field observations, and engagement in continental forums. This reflects a commitment to knowledge systems often excluded from dominant epistemic regimes, and a methodological orientation grounded in relational ethics. The data corpus is designed to center African voices, concepts, and practices as authoritative sources of ethical reasoning, in alignment with the epistemological values of Ubuntu.

A key pillar of this dissertation is the set of 26 expert interviews conducted with stakeholders across the African continent. These semi-structured interviews offer invaluable insights into the design, critique, and contextual adaptation of AI governance mechanisms. Interviews were conducted in English and French, across a mix of virtual and in-person formats. Among those interviewed, several names stand out for their contributions to critical AI discourse and practice in Africa.

Dr. Melody Musoni, former data protection advisor at the Southern African Development Community (SADC) and current analyst at ECDPM, offered a foundational critique of symbolic ethics. She observed that Ubuntu, while rhetorically invoked in many policy settings, remains ineffective without being matched with infrastructural power and legal teeth.

Benjamin Akinmoyeje, a PhD scholar from the Namibia University of Science and Technology, raised concerns about the uncritical transplantation of European ethical templates. He emphasized that “many African strategies are revised blueprints from elsewhere, filled with borrowed terminology that does not resonate with local cultures.” His perspective adds weight to the argument that ethical frameworks should not be externally imposed but co-designed with reference to African philosophical systems.

In a high-level exchange with Professor Souleymane Bachir Diagne, the Senegalese philosopher known for his work on pluralism and postcolonial ethics, we explored Ubuntu as a universalist ethic that invites the world to “make humanity together.” Diagne emphasized that Ubuntu is not merely a regional concept but a global philosophical resource capable of reforming the technocratic imagination behind AI systems. His insights are foundational to this research's normative framing.

In addition, the FieldFisher Africa Week podcast featured Kweya Chechi, Microsoft's Government Affairs and Legal Director for Africa, who provided detailed reflections on responsible AI from a policy and implementation perspective. She spoke of Microsoft's

work with the Kenyan Red Cross, AI for Good Lab, and efforts to extend broadband and AI training across underserved communities in Africa.

Strive Masiyiwa, founder and chairman of Econet, contributed views during a conversation on infrastructure, AI ownership, and innovation in Africa. He warned that “Africa must not only use AI but participate in building and owning it. Otherwise, we will be spectators in the next century.” His concerns about digital colonialism and the necessity of youth empowerment through AI governance are cited throughout this work as critical to the thesis of decolonial ethics.

In terms of events and field-based data generation, this research is informed by direct participation in major regional and global forums. These include the Africa AI Summit 2025 in Kampala, Uganda, the Global AI Summit on Africa in Rwanda, the Paris AI Summit, the Digital Rights and Inclusion Forum (DRIF) in Lusaka, Zambia, the Global Internet Governance Forum (IGF) in Riyadh, and the Africa IGF in Tanzania. At these events, insights were gathered through formal panel observations, closed-door policy labs, and informal exchanges with government officials, academics, experts, civil society actors, and private sector representatives.

Archival and documentary sources also play a critical role. This includes policy documents such as the African Union’s AI Strategy, the Data Policy Framework, and national strategies from Kenya, Rwanda, Ghana, Nigeria, and South Africa. These are not treated as fixed texts but as evolving political instruments embedded in global-local power dynamics. The Ghanaian strategy invokes “responsibility” and “trust,” but omits specific reference to indigenous ethics. South Africa’s AI framework includes “inclusion” as a value, but offers little philosophical grounding. These gaps inform the dissertation’s critique of ethical minimalism.

International frameworks such as the UNESCO Recommendation on the Ethics of AI (2021), the OECD AI Principles, and the European Commission’s AI Act have also been analyzed, particularly in relation to how they circulate in African contexts. UNESCO’s Africa Needs Assessment Report (2022) highlighted major gaps in local capacity and called for more context-sensitive frameworks. However, its emphasis remained largely on individual rights and procedural norms, rather than relational values or linguistic plurality.

The dissertation also draws from grey literature, including policy briefs, regional assessments, and workshop reports published by organizations such as Research ICT Africa, Smart Africa, the African Union Development Agency (AUDA-NEPAD), the United Nations Economic Commission for Africa (UNECA), the Internet Society, the African Declaration on Internet Rights and Freedoms Coalition, and UNESCO’s Broadband Commission. Regional civil society organizations (CSOs) like the Collaboration on International ICT Policy for East and Southern Africa (CIPESA), Paradigm Initiative provide insights into community-led ethical frameworks and digital inclusion efforts across the continent. Initiatives such as the African Observatory on Responsible Artificial Intelligence, AI4D Africa, and the Global Index on Responsible AI contribute valuable research and policy analysis on the development and governance of AI

technologies in Africa. These diverse sources provide grassroots perspectives and demonstrate ongoing resistance to extractive AI models.

Finally, this research is grounded in continuous engagement with African philosophical and historical scholarship. It draws on Ifeanyi Menkiti's idea that personhood is socially earned, not given by default, which challenges liberal notions of autonomous identity (Menkiti, 1984). Kwasi Wiredu's reflections on consensual democracy and the dangers of uncritical conceptual assimilation shape the methodological approach (Wiredu, 1996). Mogobe Ramose's Ubuntu philosophy provides an ontological and normative framework of becoming, justice, and relational ethics (Ramose, 2002). The work of Yvonne Mokgoro (1997) demonstrates the practical application of Ubuntu within constitutional jurisprudence, while Nkiru Nzegwu (2006) reminds us that any communal ethics must be attentive to gender and power.

This study also draws from older African intellectual traditions. The ethical instructions of Ptahhotep from ancient Egypt stress intergenerational wisdom and moral restraint, offering one of the earliest codified African ethics. Cheikh Anta Diop's interdisciplinary historiography is central to recovering Africa's intellectual agency and countering colonial epistemologies (Diop, 1981). Leopold Senghor's synthesis of reason and emotion in African humanism, alongside his articulation of Negritude, offers cultural insight into dignity and affective solidarity. Souleymane Bachir Diagne (2018) provides a philosophical framework for epistemic translation, grounding African ethics within global conversations. Henry Odera Oruka's (1990) work on sage philosophy is also relevant, especially his defense of African rationality beyond oral traditions.

4. Reflexivity and positionality

This dissertation emerges from a distinct epistemic and political standpoint that can best be described as Afro-optimist and Afro-communitarian. It is shaped by a conviction that Africa has not only the right but also the capacity to propose ethical frameworks and conceptual tools that speak to global challenges such as artificial intelligence. This belief in Africa's intellectual agency is neither romantic nor naïve. It is grounded in personal engagements, institutional memory, and scholarly struggle.

Afro-optimism, in this context, refers to a determined refusal to see Africa solely through the lens of deficit or dependency. It affirms that African languages, philosophies, and social imaginaries can generate normative frameworks that are not merely reactive but aspirational. This disposition informs the decision to centre Ubuntu as more than a cultural curiosity. Ubuntu is presented as a living epistemology, capable of reshaping global conversations on technology, justice, and human dignity.

As a West African researcher actively participating in the shaping of regional digital policies and continental strategies, I have moved between spaces of grassroots engagement and high-level negotiations. I have coordinated multi-stakeholder sessions, drafted policy briefs, and trained young professionals across the ECOWAS region. Yet in each of these settings, I have also encountered the subtle mechanisms of exclusion. Even where African

voices are invited, the terms of discussion often remain foreign. This structural asymmetry has motivated me to write from within, but also against, these spaces.

This research does not claim neutrality. It is consciously situated. It speaks from a place of partiality in favour of epistemic justice. The positionality assumed here rejects the illusion of detached objectivity and embraces the legitimacy of engaged scholarship. My interviews with experts across the continent were not data collection exercises in the narrow sense. They were moments of mutual clarification, of shared intellectual labour, and of affirmation that African thought must be central, not peripheral, in discussions of AI governance.

Afro-optimism is not the same as utopianism. It acknowledges the infrastructural, institutional, and epistemic challenges that shape African engagement with AI. Yet it insists that the narrative of catch-up is inadequate. Africa does not need to follow existing blueprints. It can offer alternatives. Afro-communitarian ethics, grounded in personhood, interdependence, and harmony, challenge the logic of hyper-individualism that underpins many global frameworks. These ethics are not abstract. They are lived realities, transmitted in stories, rituals, and communal deliberations.

Being positioned within multiple networks continental CSO coalitions, regional policy bodies, global fora has allowed me to observe how African values are often acknowledged rhetorically but excluded structurally. I observed recurring patterns: the ethical framing revolved around risk mitigation, transparency, and individual rights. These are important themes, yet they leave out questions of relational accountability, spiritual harm, and intergenerational responsibility. When I raised these in closed-door meetings, they were welcomed, but rarely integrated.

This disjunction between consultation and conceptual influence reveals why positionality matters. To write from Africa is not simply a geographic act. It is an epistemological and political one. It is to take seriously the notion that ethical reasoning is embedded in history, language, and community. It is to insist that African frameworks deserve to be heard not because they are exotic, but because they are insightful.

Reflexivity requires that I account for my own institutional affiliations and limitations. As someone who also operates in donor-influenced spaces, I am aware of the pressures to align language with funders' expectations. This creates moments of tension: between philosophical integrity and policy pragmatism, between what can be said and what will be heard. These tensions are acknowledged, not erased. They are part of the methodological fabric of this research.

This work is also influenced by my linguistic background and cultural worldview. My understanding of ethics has been shaped as much by traditional elders and oral teachings as by academic texts. Concepts such as “omoluabi” in Yoruba, “utu” in Swahili, or “ubuntu” in Nguni languages are not theoretical devices. They are moral compasses. They shape how people live, how they decide, and how they heal. This immersion allows me to write not only *about* Ubuntu but *from* Ubuntu.

5. Limitations

Every research undertaking carries a constellation of constraints shaped by its methodology, philosophical scope, and available sources. This dissertation, grounded in a normative and critical philosophical inquiry, seeks to enrich the ethical discourse on artificial intelligence by engaging African thought, particularly Ubuntu. However, this commitment to depth in conceptual analysis also marks the first limitation: the study does not employ empirical tools such as large-scale surveys or ethnographic immersion. As a result, the experiences of diverse stakeholders at the grassroots level are indirectly represented rather than directly observed (Crawford, 2021).

Second, the study engages a selected group of African philosophical traditions, including Ubuntu, Bantu metaphysics, Yoruba epistemology, and Akan psychological systems. These are justified by both their relevance to the theme and the availability of primary and secondary literature. However, these traditions do not exhaust the breadth of Africa's philosophical landscape. Intellectual frameworks from North Africa, the Sahelian zones, or Madagascar remain largely unexplored. This limitation arises both from the researcher's positional scope and from the lack of systematically translated and peer-reviewed literature on certain traditions (Wiredu, 1996; Masolo, 2010).

Moreover, the oral nature of many African epistemologies presents a challenge for academic treatment. As noted in African philosophical literature, orality does not preclude rigor, but its transmission structures differ fundamentally from written traditions and thus resist standard textual citation methods (Coetzee & Roux, 2002). This resistance poses a methodological constraint in framing arguments within an academic institution that prioritizes print-based scholarship.

A further limitation concerns linguistic reach. Although the research makes use of Francophone and Anglophone literature and policy documentation, perspectives from Lusophone and Arabophone Africa are notably underrepresented. This results not only from the colonial history of language hierarchies in academia but also from the researcher's own linguistic competencies and the distribution of available sources (Mekuria, 2022). The gap is partially mitigated through engagement with regional multilateral institutions such as AUDA-NEPAD, the African Commission on Human and Peoples' Rights, and pan-African CSOs, but disparities remain.

While interviews with experts and actors from across Africa offer critical insight, the sample leans towards urban, academic, and policy-centric voices. Participants include Souleymane Bachir Diagne (Columbia University), Melody Musoni (ECDPM), Benjamin Akinmoyeje (Namibia University of Science and Technology), and Kweya Chechi (Microsoft Africa), among others. Their views represent engaged thought leadership, yet the absence of perspectives from traditional leaders, rural farmers, spiritual custodians, and youth in informal sectors reflects the wider inequalities of access in AI policy consultation processes (Musoni, 2024; Diagne, 2025).

Another limitation arises from the selective engagement with international normative frameworks. This research focuses extensively on the UNESCO Recommendation on the

Ethics of Artificial Intelligence, the OECD AI Principles, and the EU's AI Act. These were chosen for their prominence in multilateral discourse and their impact on African digital policy alignment. However, frameworks emerging from China, India, and Latin America are not analyzed in equivalent detail. This narrows the comparative lens and may overlook strategic South-South ethical exchanges relevant to Africa's positioning in global technology governance (Floridi & Cows, 2019).

Temporal limitations also affect the scope of this study. Artificial intelligence is evolving rapidly, with new strategies emerging across the continent. Since the commencement of this thesis, Senegal, Ghana, Mauritius, and Rwanda have advanced national AI policies. The African Union's Continental AI Strategy is also undergoing refinements. While this research captures foundational philosophical and institutional dynamics, it cannot fully incorporate developments occurring after the finalization of its primary data and literature (UNESCO, 2022; Smart Africa, 2023).

Furthermore, the philosophical orientation of the thesis rooted in Afro-communitarianism and decolonial critique, offers both insight and constraint. The decision to affirm African intellectual sovereignty and ethical pluralism enables a radical reimagining of AI ethics, but it may under-engage more skeptical or dissenting African voices. Critical analysis of elite capture, uneven digitization, or AI-driven surveillance regimes in Africa are acknowledged but not deeply explored. Scholars such as Eubanks (2018) and Taylor (2019) highlight how AI systems may reinforce inequality even when grounded in well-meaning ethical frameworks. This tension is recognized as a limitation of scope rather than of significance.

Reflexivity introduces another constraint. As a researcher deeply embedded in digital rights advocacy, multistakeholder forums, and continental working groups, there is a risk of interpretive bias. The framing of research questions, thematic emphases, and conceptual lenses reflect lived positionality.. As scholars of critical methodology note, standpoint epistemology is both an asset and a limit (Harding, 2004; Alcoff, 2006).

Lastly, the reliance on grey literature from institutions like Research ICT Africa, CIPESA, Paradigm Initiative, AUDA-NEPAD, Smart Africa, and the Internet Society provided valuable policy insight but sometimes lacked the analytical rigor and peer-review vetting found in academic sources. While such reports are essential in understudied fields like AI in Africa, they should be complemented in future research by longitudinal studies, national archives, and community-based data.

In conclusion, these limitations are not incidental. They reflect the very terrain of African engagement with AI: fragmented, polyphonic, underrepresented, and yet rich with philosophical and political potential. This dissertation acknowledges these limits transparently and encourages future researchers to deepen, expand, and diversify the work of African-centered ethical inquiry in artificial intelligence governance.

Chapter 4: Ubuntu and AI Ethics

1. Ubuntu as relational personhood

The African philosophy of Ubuntu embodied in expressions such as "*Umuntu ngumuntu ngabantu*" (a person is a person through other persons) stands in radical contrast to dominant Western metaphysical assumptions of individualism. In Ubuntu, personhood is not a static or biologically inherited trait but a moral and social achievement attained through one's relations with others and the community. This conceptualization of the self as fundamentally interdependent provides a transformative entry point for articulating AI ethics in Africa. It challenges frameworks rooted in Cartesian dualism or liberal humanism and offers a collective, relational, and participatory ethos of dignity and moral consideration.

As Ifeanyi Menkiti (1984) famously argued, African thought prioritizes the community as the primary locus of moral value, whereas Western frameworks often center the individual as the bearer of rights and rationality. Menkiti distinguishes between the "minimal" Western conception of personhood, which is inherent at birth, and the "maximal" African conception, which must be earned through ethical conduct and social contribution. Similarly, Mogobe Ramose (2002) explains that Ubuntu places the individual within the community as an ontological necessity; there is no isolated being in Ubuntu; only being-with-others. Ramose's formulation is not merely moral but metaphysical, asserting that the relationality at the core of Ubuntu is a principle of reality itself.

This view of personhood compels a rethinking of how artificial intelligence systems interact with and reflect African societies. If AI systems are designed without regard to this ontological orientation, they risk reinforcing values that are incompatible with African moral worlds. For example, the Western model of AI personhood and agency is often infused with autonomy, control, and self-interest traits drawn from Enlightenment rationalism. Ubuntu, by contrast, locates ethical agency in intersubjectivity, mutual care, and shared flourishing. As such, AI systems deployed in African contexts should not merely operate with technical precision but must align with communal values, moral accountability, and restorative outcomes.

Yvonne Mokgoro, a former justice of South Africa's Constitutional Court, used Ubuntu to argue for jurisprudence centered on reconciliation, compassion, and dignity. She saw Ubuntu not as a soft moral code but as a constitutional value that grounds human rights and governance in South Africa. This approach translates well into AI ethics, where procedural justice and dignity should not be mere compliance checkboxes, but core design imperatives. AI ethics guided by Ubuntu does not begin with "what is fair for the individual?" but "what sustains the community?" The normative horizon is communal flourishing, not isolated efficiency.

From a governance standpoint, Ubuntu also suggests a participatory model. Kwasi Wiredu (1996) argued for a consensual democracy rooted in African traditions of deliberation, where decisions are made through dialogue rather than majoritarian imposition. This model can inform multistakeholder AI governance in Africa, which must move beyond token inclusion to meaningful participatory design and deliberation. In practice, this could translate into community-based AI policy consultations, culturally anchored data

governance models, and ethical frameworks that reflect both the lived experiences and aspirations of African societies.

Beyond ethics, Ubuntu contributes to epistemology. Nkiru Nzegwu emphasizes that African epistemologies are dynamic, performative, and relational knowledge is not merely what is observed, but what is practiced and shared. This view challenges the data-extractivist logic underlying many AI systems, where knowledge is reduced to patterns and probabilities derived from unethically gathered data. Ubuntu foregrounds consent, purpose, and communal benefit, offering a reparative framework for data practices in Africa. An Ubuntu-informed AI design must center the values of transparency, reciprocity, and contextual relevance.

Reflecting on my own research journey, these insights were not simply theoretical. In conducting expert interviews, including conversations with thinkers like Souleymane Bachir Diagne, Melody Musoni, Strive Masiyiwa, Benjamin Akinmoyeje, and others, it became evident that Ubuntu is more than a South African ideal; it is a continental moral grammar. Diagne, for instance, underscored Ubuntu's potential to "make humanity together," highlighting its role in resisting ethno-nationalism and environmental degradation. Musoni warned of the dangers of reducing Ubuntu to rhetoric if not matched with investment, infrastructure, and institutional reform. These voices confirmed the urgency of grounding AI governance in living traditions rather than imported templates.

This view was echoed at the Global AI Summit in Rwanda, where panelists repeatedly called for ethical models that reflect Africa's historical and philosophical realities. Ubuntu is not a nostalgic return to precolonial purity but a philosophical tool for modern governance: flexible, dialogical, and responsive to complexity. It offers a counterpoint to the abstraction and universality of dominant AI ethics frameworks, instead grounding normative claims in the relationships, obligations, and moral duties that sustain African life.

Ubuntu also resonates with decolonial theories that critique the Eurocentric foundations of technological development. Scholars like Achille Mbembe, Walter D. Mignolo, and Sylvia Wynter have warned against the "coloniality of power" that persists in global AI systems. Ubuntu resists this hegemony by proposing a local ethical anchor, one that is neither exclusionary nor parochial. As I argued during my lecture on "Relationality and Humanization vs Alienation," Ubuntu is a site of both resistance and creativity. It humanizes technology by binding it to the rhythms of African life.

As AI increasingly mediates identity, knowledge, and governance, Ubuntu challenges us to ask: What kind of persons are we creating in our systems? What values do these systems amplify or erase? In Ubuntu, to be human is to be accountable to others. This accountability must be programmed into the logics of AI systems deployed in Africa. It is not enough to regulate AI, we must also reimagine it through the prism of relational personhood.

2. Implications for data identity

If Ubuntu redefines personhood in relational terms, it also invites a reconceptualization of data identity. In dominant AI systems, data is abstracted, decontextualized, and stripped of its social roots. This transformation allows it to be stored, categorized, analyzed, and traded across global platforms. However, in an Ubuntu-centered ethical vision, data cannot be seen as a neutral byproduct of individual activity. It is a moral extension of the person and, by implication, the community. Understanding this has profound implications for how African societies design, govern, and resist AI systems.

In most Western AI frameworks, personal data is treated as property or as a transactional asset subject to consent-based agreements. This model assumes that individuals are autonomous agents capable of understanding and negotiating complex terms of use. Yet as highlighted in the interviews with Dr. Melody Musoni and Benjamin Akinmoyeje, this vision does not reflect the reality on the ground in many African contexts, where data about an individual often implicates a broader social network family, village, or ethnic group. Dr. Musoni warned against a superficial reading of Ubuntu that does not come with infrastructure, investment, and legal innovation. She emphasized that “relational personhood also means relational rights and relational risks,” and that this must be reflected in how data is gathered and processed.

This view is supported by literature on data justice, which argues that data rights cannot be limited to individual autonomy. According to Taylor (2017), data justice involves understanding how power is exercised through data, how it affects marginalized communities, and how collective agency can be reasserted. Ubuntu, in this sense, offers a compelling moral and political grounding for African data justice movements. It shifts the focus from individual consent to communal accountability and from ownership to stewardship.

In practice, this means that the ethical use of data must account for context, ancestry, and locality. Take, for example, the deployment of AI in healthcare. In many African countries, biometric data is collected through national identification programs, often funded by international donors or private tech firms. These datasets are sometimes shared across ministries or with third parties without community consultation. From a Ubuntu perspective, this constitutes not only a breach of privacy but also a rupture of social trust. The face, fingerprint, or genome of a person is not just a technical input, it is a symbol of belonging, of shared fate. Mishandling it damages more than the individual; it weakens the moral fabric of the community.

In the Yoruba worldview, the notion of *omoluabi* denotes a person of character who lives in harmony with others. This is echoed in Swahili with *utu*, in Bantu languages with *botho*, and in many other African moral vocabularies. These concepts are not only ethical ideals but also frameworks for how knowledge, including digital knowledge, should be respected and circulated. For instance, to gather data about agricultural practices from a rural community is not just to extract information, it is to engage in a relationship of trust, where the return of benefit and recognition is expected.

African philosophers such as Kwasi Wiredu and Nkiru Nzegwu have long argued that traditional African knowledge systems are rooted in consensus, negotiation, and reciprocity. These values are poorly reflected in algorithmic logics that prioritize efficiency and prediction over meaning and consent. In his analysis of consensus, Wiredu (1996) emphasizes that truth in African epistemology is not adversarial or objective in the Western sense, but socially validated through ongoing dialogue. When applied to data systems, this would require collective deliberation on what counts as legitimate data use, who gets to decide, and how communities can challenge unjust deployments.

One illustrative case comes from the AI4D Africa Network's 2023 report on AI in agriculture. In one East African pilot, drone surveillance was introduced to monitor crop cycles and predict yield. While the data improved logistical planning, local farmers reported discomfort with the constant monitoring of their fields, which clashed with customary land rights and seasonal rituals. The ethical guidelines had included informed consent, but only at the individual level. No traditional authorities or collective landholders were consulted. This reflects what scholars like Milan and Treré (2019) call “epistemic extractivism”, the appropriation of local knowledge and data without reciprocal benefit or recognition.

The African Union's Data Policy Framework, adopted in 2022, begins to address these gaps by calling for people-centered approaches to data governance that align with African values. However, its implementation remains fragmented. As one participant in the 2025 Africa AI Summit in Kampala noted, “Policies on paper are not enough. Without participatory mechanisms and local oversight, AI will repeat the same mistakes of top-down development.” This sentiment is echoed in grassroots work done by the Internet Society and Research ICT Africa, which call for decentralizing data governance and supporting digital literacy initiatives rooted in community languages and priorities.

Ubuntu also questions the metaphysical assumptions of AI systems. Data is often treated as raw material value-neutral, inert, and waiting to be refined. But Ubuntu understands knowledge as a living force, embedded in relationships and infused with spiritual meaning. The data of a person is not only about them; it is part of who they are. Misuse of that data is not just unethical; it is a form of violation, a disrespect for personhood. This can be seen in indigenous conceptions of naming, memory, and lineage, where identity is co-produced across generations.

This ontological insight matters when we consider technologies like facial recognition, language models, and digital ID systems. If these systems treat people as isolated data subjects, they undermine the very moral architecture of Ubuntu societies. A language model trained exclusively on English or French fails to capture the moral grammars of *Ewe*, *Hausa*, *Amharic*, or *Fula*. The exclusion is not merely linguistic, it is ethical.

To move forward, African states and institutions must embed Ubuntu in data governance infrastructures. This includes not only laws and regulations, but also practices: community ethics councils, oral consent rituals, culturally sensitive interfaces, and mechanisms for collective redress. Partnerships with grassroots movements, traditional leaders, and youth innovators are essential to this process.

Ubuntu's implications for data identity are transformative. They call for a reconceptualization of data as relational, sacred, and situated. This is not just a philosophical proposition; it is a political imperative. As Africa charts its AI future, the ethics of data must be grounded not in imported norms but in the living values of the people who generate it.

3. Comparison with Western ethics

The ethical landscape that governs artificial intelligence is deeply influenced by the philosophical traditions that inform societal values, legal systems, and governance norms. When juxtaposing Ubuntu philosophy with dominant Western ethical frameworks in AI, it becomes evident that they rest on fundamentally divergent ontologies of personhood, moral reasoning, and social responsibility. Ubuntu, an African philosophical paradigm emphasizing interdependence, dignity, and communal flourishing, offers a striking contrast to the individualist, rights-based, and utilitarian foundations of many Western approaches. This divergence is not merely conceptual but deeply consequential in shaping AI policies, especially when ethical systems translate into computational logics, regulatory standards, and data governance regimes.

Ubuntu's emphasis on relational personhood summarized in the often-quoted phrase "I am because we are" displaces the notion of an atomized, autonomous individual as the moral center. Instead, it centers the community as the site of moral reasoning and human identity. This contrasts sharply with ethical theories dominant in Western AI governance discourse.

For example, deontological ethics, influenced by Kantian thought, focuses on the rights and duties of individuals, assuming a rational agent capable of moral deliberation in isolation. Similarly, consequentialist approaches like utilitarianism assess moral worth based on aggregate outcomes, often through quantifiable means such as cost-benefit analyses. Both of these frameworks, while internally coherent, are ill-equipped to account for the dense social, historical, and relational fabrics that characterize African societies.

The implications of this divergence become stark in the design and deployment of AI systems. Western models often embed assumptions of objectivity and neutrality, which obscure structural inequalities and historical injustices. Ubuntu ethics, by contrast, foregrounds historical situatedness and social harmony. It insists that technologies must promote healing, solidarity, and collective well-being. For instance, Ubuntu's approach to justice is restorative rather than retributive; this informs not only post-conflict reconciliation in countries like South Africa but also digital justice. AI systems designed with Ubuntu in mind would prioritize repairing harms caused by algorithmic bias, ensuring not only accountability but restoration within the communities affected.

Historical trajectories further illuminate these contrasts. Western ethical frameworks have evolved alongside the Enlightenment project, which valorized individual reason and universalism. This trajectory led to the codification of ethics into formal rights-based frameworks such as the Universal Declaration of Human Rights and more recently, the European Union's General Data Protection Regulation (GDPR) and the EU AI Act. These instruments offer legal precision but can fall short of addressing ethical pluralism and

contextual nuance. In African contexts, where colonial legacies have left deep scars on institutions and epistemologies, these frameworks may appear externally imposed or alienating. Ubuntu, rooted in precolonial modes of governance and dispute resolution, offers a normative anchor that is both endogenous and forward-looking.

Philosophers such as Kwame Gyekye and Mogobe Ramose have long emphasized that African ethics resist the binary between the individual and the collective. Instead, they present a layered subjectivity where personal dignity is achieved through relationships. Ramose (2002) argues that Ubuntu is both ontological and ethical, a mode of being and a way of doing. This is significant when designing AI systems that must engage with users not as isolated data points but as beings embedded in social webs. Western algorithms trained on anonymized, decontextualized data risk erasing cultural specificity. Ubuntu invites us to ask: What relationships does this AI system enhance or erode? Who is left out of its epistemic horizon?

Moreover, as Melody Musoni highlighted in her interview, Ubuntu's relationality cannot be reduced to rhetoric. Without institutional investments, infrastructure, and inclusive governance mechanisms, Ubuntu risks becoming an empty signifier. This critique echoes broader concerns within African philosophy, where scholars like Paulin Hountondji have warned against romanticizing tradition without interrogating its political economy. The call is not for a wholesale rejection of Western ethics, but for a pluralistic framework that allows for epistemic justice and cognitive justice in the digital age. This aligns with reflections from Souleymane Bachir Diagne, who advocates for an Ubuntu-inflected cosmopolitanism that embraces dialogue across traditions while resisting epistemic erasure.

In practice, Western AI ethics often emphasize principles such as transparency, explainability, and fairness. These are undoubtedly important, but when applied universally without contextual grounding, they can reproduce the very inequalities they aim to solve. Ubuntu pushes us beyond principles into praxis. How are AI systems experienced in the everyday life of rural farmers, informal traders, or schoolchildren in African settings? What communal obligations do developers, data scientists, and regulators have to these communities? The African Commission on Human and Peoples' Rights has echoed this view, insisting on AI systems that enhance human dignity and align with the needs of the continent's peoples.

These insights are not purely theoretical. Initiatives such as the African Union's High-Level Panel on Emerging Technologies, the African Research Centre on AI in Congo Brazzaville (ARCAI), and Research ICT Africa's Just AI Centre are actively working to integrate Ubuntu-inspired ethics into policy frameworks. These organizations recognize that Africa's AI governance cannot rely solely on imported templates. Instead, they advocate for participatory governance, community-based impact assessments, and ethical review boards that reflect the continent's diverse moral traditions.

The tension between Ubuntu and Western ethics is not antagonistic but dialogical. In my reflections, I have come to see this dialogue as necessary for constructing truly inclusive AI governance. Western models bring technical rigor and institutional depth, while Ubuntu offers moral imagination and a politics of care.

4. Ubuntu and Islamic Ethical Traditions in African Moral Thought

While this dissertation primarily engages Ubuntu as a philosophical framework rooted in sub-Saharan African traditions, any discussion of ethics on the African continent would be incomplete without acknowledging the influence of Islamic moral thought. Islam has shaped ethical reasoning, governance practices, and communal life across large parts of Africa for centuries, particularly in North, West, and East Africa. Rather than existing in opposition, Ubuntu and Islamic ethics have historically coexisted, interacted, and, in many contexts, mutually reinforced one another.

Islamic ethical traditions emphasise concepts such as *umma* (community), *adl* (justice), *amanah* (trust), and *maslaha* (the common good). These principles resonate strongly with Ubuntu's relational conception of personhood, collective responsibility, and moral accountability. Both ethical systems prioritise the well-being of the community over radical individualism and frame moral action as inseparable from social obligation. In many African societies, Islamic jurisprudence and customary norms informed by Ubuntu have developed in parallel, producing hybrid moral ecosystems where ethical legitimacy derives from both communal consensus and religious duty.

This convergence is particularly relevant for contemporary debates on artificial intelligence governance. Islamic ethics, like Ubuntu, approaches moral responsibility as relational and contextual rather than purely procedural. Concepts such as stewardship (*khilafah*) and trust provide ethical resources for thinking about data governance, accountability, and intergenerational responsibility in ways that complement Ubuntu-informed perspectives. Together, these traditions challenge dominant AI ethics frameworks that rely primarily on individual autonomy and contractual consent.

Chapter 5: Governance, Policy and Legal Implications of Ubuntu-Informed AI Ethics

1. Ubuntu and Precolonial African Governance Models as Ethical Foundations for AI

Artificial intelligence has catalyzed a profound ethical reckoning in our time. While global frameworks have attempted to grapple with AI's complex impact, most of these remain steeped in Euro-American paradigms, emphasizing rational individualism, utilitarian outcomes, or abstract notions of fairness. Yet, Africa is not without its own repositories of ethical reflection. Long before the formalization of data protocols and algorithmic accountability, African societies were already navigating governance challenges through relational philosophies rooted in harmony, balance, and the collective good. Among these traditions, Ubuntu emerges not only as an ethical compass but as a functional governance model that predates colonial disruption and remains relevant to modern AI dilemmas.

. Precolonial African governance systems, including those of the Akan in present-day Ghana, the Bantu-speaking communities of Southern Africa, and the Oromo Gadaa system in Ethiopia, offer a wealth of case-based governance logic centered on inclusive deliberation, moral reciprocity, and non-retributive justice. For instance, in traditional Igbo societies, decision-making was often a communal affair, mediated through councils of elders, women's assemblies (Umuada), and age-grade associations structures that ensured layered accountability and participatory consensus rather than unilateral rule.

Applied to AI, these models challenge the dominant logic of centralized control and opaque algorithmic processes. Ubuntu ethics stress transparency through collective witnessing (orality and consensus), justice through restoration rather than punishment, and dignity as a communal, not individual, asset. In this view, ethical AI must prioritize the relationships it mediates, not just the outputs it optimizes. It must also be aware of its long-term impact on intergenerational cohesion, ecological balance, and local identity, themes that are often absent from technocratic governance models.

The adoption of Ubuntu as a philosophical lens also invites a reconsideration of what constitutes a “harm.” Western bioethical models often reduce harm to quantifiable individual injury, yet Ubuntu-inspired reasoning would include communal fragmentation, alienation, and loss of dignity as equally urgent ethical concerns. For example, the deployment of biometric surveillance tools in African urban centers, often justified as enhancing security, might satisfy compliance checklists from global vendors. But when evaluated through Ubuntu, the same systems may be judged unethical if they sever bonds of trust in neighborhoods, reproduce colonial surveillance practices, or erode cultural norms of neighborly care.

Moreover, the spiritual dimension of Ubuntu should not be overlooked. Precolonial African systems often fused the sacred with the civic, acknowledging ancestors, land spirits, and collective memory as integral to legitimate decision-making. While secular AI frameworks tend to eschew such dimensions, ignoring them in the African context risks perpetuating

epistemic injustice. For many communities, ethical authority flows from continuity with the past, not just legal compliance or mathematical neutrality. This has direct implications for how data is sourced, how consent is understood, and how AI models should be governed over time.

Institutionalizing these principles will require more than rhetorical gestures. It demands the development of governance protocols that allow communities not just policymakers to define ethical thresholds and participate in audits. This could mean integrating traditional councils or faith-based forums into AI regulatory consultations, ensuring that policy design reflects not only formal legality but lived legitimacy. International frameworks like the UNESCO Recommendation on the Ethics of AI already encourage such pluralism, but implementation remains thin across Africa.

This historical grounding also underscores that Ubuntu is not a mere philosophical abstraction. It is a practiced and practiced ethic, embedded in real institutions that governed land, marriage, trade, conflict resolution, and even early forms of information management. The use of griots, storytellers, and memory custodians across the Sahel and West Africa is a testament to indigenous approaches to data stewardship approaches that valued contextual integrity, intergenerational knowledge, and collective verification long before the digital era.

Thus, positioning Ubuntu as a foundation for ethical AI is not an act of cultural nostalgia but a forward-looking strategy for building inclusive, sustainable, and context-sensitive digital futures. African states now have an opportunity to frame AI ethics not as a passive adoption of global norms but as a renaissance of indigenous wisdom renegotiated for the digital age.

2. Designing a Policy Framework Grounded in Ubuntu Values

Designing an ethical AI policy framework anchored in Ubuntu requires moving beyond rhetorical invocations and toward a structurally embedded philosophy that reshapes both the substance and process of governance. Unlike many contemporary frameworks that begin with risk assessments and end with compliance checklists, an Ubuntu-based approach starts from a different epistemological premise: that human beings are defined through relationships, and ethical governance must therefore sustain and enhance the web of relational life.

The policy architecture must be participatory, pluriversal, and situated. This implies first a radical commitment to co-creation. Communities are not to be merely consulted; they must be empowered to shape AI policy through deliberative forums, local councils, and indigenous institutions. This aligns with the traditional practice of *indaba* in Southern Africa, where decisions are reached through inclusive consensus and deep dialogue. As Metz (2011) argues, Ubuntu's procedural ethic privileges dialogical reasoning, not just outcomes, and thereby favors policies built through communal affirmation over those imposed by technocratic fiat.

The framework must also address epistemic justice. This includes incorporating concepts and languages from African knowledge systems. For instance, instead of merely referring to “privacy” in the liberal individualist sense, policymakers should consider locally resonant values like *omoluabi* in Yoruba (a person of good character accountable to the community) or *utu* in Swahili (a moral character grounded in mutuality). These terms capture a richer and more socially embedded ethical vocabulary. The African Union’s continental AI strategy encourages cultural contextualization (African Union, 2023), but national policies often default to imported standards. To remedy this, policy guidelines should mandate multilingual documentation, culturally embedded ethical training, and indigenous representation in review boards.

In terms of institutional design, a decolonial Ubuntu framework would include a Community Ethics Review Council (CERC) as a standing body, composed of traditional leaders, spiritual custodians, civil society actors, youth representatives, and technologists. This CERC would be responsible for providing ethical oversight, reviewing AI deployments, and ensuring they align with Ubuntu principles such as solidarity (*umunthu*), restorative justice, and relational accountability. Such a structure reimagines the function of AI regulators not as remote enforcers of imported norms, but as ethical stewards embedded in cultural realities.

An Ubuntu framework would also require a new mode of impact assessment. Instead of limiting itself to individual rights or statistical fairness, it would mandate what can be termed a Relational Impact Assessment (RIA). This tool would evaluate the effects of AI systems on social cohesion, generational trust, ecological harmony, and cultural continuity. For example, a chatbot introduced in a rural area for agricultural advice would be assessed not only for accuracy and bias but also for whether it enhances local knowledge practices, empowers women, or erodes traditional farmer cooperatives.

These ideas have practical precedents. In the grassroots AI translation work of the Masakhane project, African NLP practitioners have built tools to preserve and promote African languages using community co-authorship and open governance. The success of this model lies not only in its outputs but in its ethical process, a collective stewardship of digital resources that embodies Ubuntu’s ethic of participation and care (Masakhane, 2022). Similarly, Rwanda’s AI strategy has experimented with community-driven digital inclusion pilots that embed values of reconciliation and national unity, showing how ethical imperatives can be woven into state-led technological visions (Government of Rwanda, 2021).

To translate Ubuntu values into policy, legal recognition is also required. For instance, the constitutions of countries such as South Africa and Kenya already reference cultural values and customary law. An Ubuntu AI framework would advocate for explicit legislative provisions recognizing community data ownership, requiring consent not just at the individual but collective level. This aligns with reflections shared by Melody Musoni, who emphasized that data governance in Africa cannot ignore collective memory, lineage, or land ethics, which are central to communal rights and obligations (Musoni, 2025, personal interview).

The ethical principles of Ubuntu also require rethinking intellectual property. Knowledge generated from community interaction such as local health practices used to train AI diagnostics should be considered communal intellectual heritage. Hence, AI frameworks must integrate benefit-sharing protocols, echoing the Convention on Biological Diversity, but adapted to the digital realm. Here, the policy logic must resist extractivism and instead embody restitution, reciprocity, and shared stewardship.

In line with this, Strive Masiyiwa emphasized that AI innovation must benefit African communities not merely as users but as owners and co-creators. He argued for public-private-community partnerships that reinvest in digital infrastructure, education, and youth entrepreneurship to prevent a new form of digital dependency (Masiyiwa, 2025, personal interview). Policy design must therefore allocate AI funds toward communal infrastructure such as data cooperatives, rural innovation labs, and cultural archives and ensure they are governed not by donor conditionality but by local consensus.

Ubuntu-based AI governance also recognizes temporality differently. It centers intergenerational continuity, meaning that AI deployment should be evaluated not only for immediate efficiency but for its impact on future generations. Inspired by African cosmologies that honor ancestors and descendants in decision-making, AI policy should institutionalize mechanisms for intergenerational representation, perhaps through rotating youth-elder review panels that examine long-term cultural, environmental, and psychological consequences.

Moreover, while Ubuntu favors harmony, it does not mean avoiding conflict. On the contrary, conflict resolution mechanisms are central to Ubuntu traditions. Policy frameworks should therefore incorporate restorative justice mechanisms for AI-related harms. For example, when algorithmic decisions cause exclusion, communities should have access not just to litigation but to culturally sensitive mediation forums. These forums could be embedded within national AI ombudsman offices, modeled after traditional restorative councils.

Designing such a policy framework is not merely about ethics; it is about sovereignty. As Benjamin Akinmoyeje emphasized in his interview, African countries cannot afford to be recipients of AI governance templates. They must build from their own traditions, institutions, and philosophical resources to offer the world an alternative model of digital governance grounded in dignity, relationality, and justice (Akinmoyeje, 2025, personal interview).

Thus, an Ubuntu-informed framework is not a moral supplement to technical regulation; it is a comprehensive reordering of what counts as knowledge, who counts as a stakeholder, and how justice is envisioned. It challenges African policymakers to resist mimetic policy transfer and instead draw from ancestral models, lived traditions, and communal visions to govern AI in a way that affirms African futures.

3. Implementation Challenges and Tensions Between Indigenous Values and Modern Regulatory Paradigms

While the ambition to develop Ubuntu-informed AI policy frameworks holds promise for ethical innovation and epistemic justice, its practical implementation faces several complex challenges. These challenges lie at the intersection of epistemological incompatibility, institutional inertia, global economic asymmetries, and internal contradictions within postcolonial African states.

One of the most persistent tensions is the epistemological divergence between relational ethics, as embodied in Ubuntu, and the logic of most regulatory frameworks, which are grounded in Western liberal legal traditions. As noted by Ramose (2002), Ubuntu cannot be fully captured through the individualist constructs of rights and responsibilities that dominate constitutional and regulatory design across Africa today. Ubuntu operates through norms of collective obligation, moral entanglement, and intergenerational memory, features that resist translation into positivist regulatory forms such as statutory law or performance metrics. This disjuncture produces friction when governments attempt to graft Ubuntu onto technocratic processes without addressing their deeper philosophical roots.

Regulatory structures often demand codification, enforcement mechanisms, and measurable benchmarks. Ubuntu, however, thrives in context-dependent application, oral reasoning, and moral persuasion. As Wiredu (1996) warned, the pressure to institutionalize African thought using Western instruments risks distorting its meaning. Efforts to implement Ubuntu in AI governance run the risk of superficiality treating Ubuntu as a slogan rather than a system of thought capable of guiding systemic regulation. This risk was echoed in the Fieldfisher podcast with Kweya Chechi, who cautioned that without adequate infrastructural investment and policy space, Africa's ethical aspirations may remain decorative (Chechi, 2025, podcast interview).

In interviews conducted for this study, Melody Musoni emphasized that policymakers often use Ubuntu symbolically to demonstrate cultural relevance without operationalizing it. She noted that many government AI strategies refer to “inclusivity” or “cultural grounding” but proceed to adopt the EU AI Act or OECD principles verbatim. This mimicry, she argued, undermines policy sovereignty and fails to address the unique ethical stakes in African contexts (Musoni, 2025, personal interview). Her observations align with the critique of “empty decolonialism” offered by Tuck and Yang (2012), where structural colonial continuities are preserved under a veneer of local participation.

There is also the challenge of navigating Africa's legal pluralism. Ubuntu-based ethics coexists with other legal and moral systems: Islamic jurisprudence in the Sahel, customary law in rural polities, and constitutional liberalism in urbanized centers. These coexisting regimes create tension in AI governance. For example, while Ubuntu emphasizes communal restitution, constitutional courts may require individual compensation. Likewise, regulatory sandboxes piloted in Smart Africa initiatives often rely on Western risk matrices, sidelining African ontologies of justice, spirituality, or the sacredness of life. Harmonizing these competing frameworks without subordinating Ubuntu to technical rationality remains a formidable challenge.

Economic dependency further complicates implementation. As Benjamin Akinmoyeje noted in his interview, many African states are constrained by external donor requirements and international investment frameworks that prioritize compliance with global standards over local innovation. He warned that frameworks shaped primarily to attract FDI or align with World Bank guidelines may marginalize community-led governance or alternative knowledge systems. Indeed, African digital policy is often caught between the imperatives of modernization and the ethics of rootedness. The desire to be globally competitive can lead to the adoption of AI ethics models that prioritize efficiency, accountability, and risk management over solidarity, reciprocity, and communal well-being (Akinmoyeje, 2025, personal interview).

Institutional capacity is another key barrier. Many African countries lack national AI ethics commissions, regulatory agencies, or digital rights ombudsman offices. Where such structures exist, they are often underfunded and understaffed. Embedding Ubuntu into AI governance requires time, skilled personnel, and dedicated funding—not only for policymaking but for community consultations, local-language translation, traditional leader engagement, and iterative feedback loops. Without long-term investment, the vision of Ubuntu AI governance risks being reduced to a few paragraphs in a national strategy document, rather than embedded into practice.

Moreover, internal contradictions within postcolonial African states cannot be ignored. Authoritarian tendencies, politicization of data governance, and the militarization of digital infrastructure can erode trust and undermine communal values. For instance, some governments promote Ubuntu while simultaneously deploying surveillance technologies that threaten collective freedoms. The 2024 Londa report on Togo highlighted such contradictions, noting how principles of participation and ethical oversight were weakened by security-driven logics (Paradigm Initiative, 2024). These contradictions must be addressed if Ubuntu is to be more than aspirational.

The diversity of African traditions also presents a philosophical challenge. Ubuntu is not a monolith. While its ethical grammar centered on communal interdependence and relational accountability is widely shared, its articulation differs across linguistic and cultural zones. Concepts such as *bomoto*, *omoluwabi*, *utu*, or *bodiko* reflect variations of Ubuntu in Zulu, Yoruba, Swahili, and Shona contexts, respectively. Policymakers must resist the temptation to standardize Ubuntu into a universal checklist and instead allow for regional and linguistic plurality. This plurality should be seen as a strength, enabling context-sensitive policies that are resonant with local moral imaginaries.

Finally, there is a danger of epistemic tokenism. As Fricker (2007) reminds us, epistemic injustice occurs not only through exclusion but also through superficial inclusion that instrumentalizes knowledge. If African philosophies are included only to legitimize pre-determined policies or to fulfill diversity mandates, they are not truly transformative. Authentic implementation of Ubuntu ethics requires epistemic humility, a willingness to learn from African traditions, to be reshaped by them, and to let them inform the very categories through which we define fairness, consent, or harm.

The way forward in embedding Ubuntu into AI governance requires a comprehensive and context-sensitive approach. First, it involves the development of African AI regulatory sandboxes that are not merely technological testing grounds but are explicitly informed by customary law. These sandboxes must be grounded in legal pluralism, recognizing that traditional authorities, local norms, and community protocols possess legitimacy equal to statutory codes. Such models could allow for experimentation with AI systems in ways that respect land-based ethics, oral consent mechanisms, and local rituals of deliberation, thereby offering frameworks that reflect African epistemologies rather than importing risk-based paradigms from the Global North.

Second, national forums must be convened that bring together elders, youth, and community leaders to deliberate on the ethical dimensions of digital technologies. These forums would serve as bridges between generations and knowledge systems, enabling the co-creation of normative guidelines that reflect lived experiences and intergenerational wisdom. In these gatherings, ancestral memory, communal accountability, and the social meaning of technology would be debated in local languages, reframing AI not as a purely technical issue but as a societal and moral question. Such an approach would help ground AI governance in Africa's traditions of consensus and participatory dialogue.

Third, there is a pressing need for curriculum reforms across African universities and training institutions. Ethical AI education should include indigenous digital philosophies and African political thought, taught not as exotic supplements but as foundational theories. These teachings must be situated alongside global ethics frameworks to foster epistemic pluralism. By introducing thinkers such as Kwasi Wiredu, Ifeanyi Menkiti, and Nkiru Nzegwu into computer science and public policy syllabi, students will be trained to navigate both the universalist aspirations of global AI standards and the situated knowledge systems that shape African societies. This pedagogical shift is essential if future technologists, regulators, and scholars are to develop frameworks that are at once innovative, inclusive, and culturally coherent.

Institutions such as the African Commission on Human and Peoples' Rights (ACHPR), ECOWAS, and the African Union's High-Level Panel on Emerging Technologies (APET) have begun opening spaces for such conversations. But their frameworks must move from consultative documents to enforceable mechanisms.

4. Legal Responses to Critiques of Ethno-Philosophy in African AI Governance

The criticism of ethno-philosophy has long shaped debates around the legitimacy of using indigenous African knowledge systems in formal epistemology, jurisprudence, and governance design. Within the landscape of AI policy, this tension is particularly acute. Skeptics argue that frameworks like Ubuntu, derived from oral traditions and communal epistemologies, lack the logical rigor, universality, and institutional specificity needed for legal codification. Yet this critique often rests on Eurocentric assumptions that privilege written text and individualist rationality as the only credible forms of knowledge production. In this section, I explore counterarguments to such critiques, drawing upon African philosophical thinkers who have challenged the marginalization of indigenous

thought and have articulated the legal and normative value of communitarian African ethics.

One of the early critiques came from Paulin Hountondji, who cautioned against the uncritical celebration of traditional wisdom without philosophical interrogation. However, subsequent generations of scholars have gone beyond his critique, developing rigorous methodologies for engaging indigenous thought systems. Kwasi Wiredu, for example, emphasized the possibility of conceptual decolonization through a dialogic process that reinterprets indigenous ideas in modern contexts. Wiredu did not reject tradition outright; rather, he proposed “conceptual translation,” which involves expressing communal ideas in rational formats compatible with contemporary policy frameworks. This is especially relevant for AI governance, where the collective orientation of Ubuntu can be rearticulated through participatory policymaking, impact assessments grounded in social values, and relational models of consent and data justice.

The Yoruba philosopher Sophie Oluwole went further by arguing that African thought is not merely a set of folk practices but a coherent system of knowledge production rooted in logic and metaphysics. Her work on Ifá divination, for instance, demonstrates the presence of probabilistic reasoning and ethical deliberation in indigenous systems that are often mischaracterized as mystical. Translating her insights to the AI field offers a rich vein of potential for rethinking probabilistic algorithms not merely in terms of optimization but as tools that require ethical calibration rooted in social accountability. Rather than treating Ubuntu as a poetic metaphor, Oluwole’s work invites us to see it as an operational logic one that emphasizes deliberation, consensus, and restorative correction rather than punitive determinism.

In a similar spirit, Mogobe Ramose asserts that law in the African context must be understood through the lens of Ubuntu ontology. His formulation of “law as wound and healing” reframes justice from being about retribution or efficiency to being about restoring harmony. This has profound implications for AI systems that today are primarily designed around risk, efficiency, and deterrence. Ramose’s reinterpretation of African legal thought offers a counterweight to techno-legal formalism, opening pathways for algorithmic systems to be embedded with values of relationality and care. These perspectives reveal how what is often dismissed as “ethno-philosophy” holds the potential to enrich contemporary legal theory and improve the ethical robustness of AI systems.

Another key contributor to this debate is Nigerian sociologist and feminist theorist Ifi Amadiume, who urges us to resist colonial binaries of public/private, male/female, and state/community in law. Her analyses of matriarchal legal forms and flexible gender structures challenge Western juridical norms and introduce the possibility of more fluid, inclusive regulatory paradigms. AI systems that replicate patriarchal or heteronormative assumptions can thus be ethically interrogated through African feminist jurisprudence. By incorporating her insights, one can envision a form of AI governance that is intersectional, historically informed, and locally grounded.

These thinkers collectively rebut the charge that African indigenous thought lacks normative weight or conceptual clarity. They offer us strategies not just to defend ethno-

philosophy, but to elevate it as a legitimate and necessary foundation for policy design in the age of AI. Rejecting ethno-philosophy outright, as some critics have done, amounts to epistemic injustice, silencing a world of ideas that emerged through millennia of communal living, oral memory, and environmental stewardship.

From a personal standpoint, embracing these philosophies in my own research has involved unlearning the disciplinary biases I was trained in and revaluing traditions I once considered “non-academic.” It was humbling to see how narratives transmitted through proverbs, folktales, and rituals often carried ethical complexity more relevant to today’s AI dilemmas than many formal policy documents. The emphasis on reconciliation, intergenerational justice, and human dignity that permeates these traditions can and should inform legal models moving forward. As Africa builds its digital sovereignty, it must also reclaim its epistemic sovereignty.

Therefore, any critique of Ubuntu or ethno-philosophy as “pre-modern” or “insufficiently legal” fails to grasp the dynamic evolution of these traditions and their capacity for integration into modern regulatory structures. Ubuntu, when properly understood through voices like Ramose, Oluwole, Amadiume, and Wiredu, offers more than moral inspiration

5. Embedding Ubuntu Principles into Formal Legal Frameworks

Translating Ubuntu into actionable legal frameworks requires a deliberate engagement with history, jurisprudence, and institutional structures. Ubuntu is not an abstract philosophy confined to oral tradition; it has historically informed dispute resolution mechanisms, customary courts, and public authority across many African societies. The task before African policymakers and legal scholars is not to “invent” Ubuntu as a legal tool, but to reassert it as a foundational paradigm that was displaced during colonization. Through this reassertion, Ubuntu can guide ethical governance in emerging domains like artificial intelligence.

One of the clearest historical precedents for Ubuntu-informed governance lies in traditional justice systems such as the Gacaca courts in Rwanda, the Bashingantahe in Burundi, and the Kgotla in Botswana. These mechanisms emphasized reconciliation, public deliberation, and community restitution over punitive justice. As Ramose (2001) has argued, these institutions were more than a conflict resolution forum, they were enactments of Ubuntu in legal form. In post-conflict societies, they served to restore social harmony rather than reinforce legalistic individualism. Integrating such models into AI governance would require that decision-making processes surrounding data, algorithms, and automation prioritize the restoration of communal trust and not just individual redress.

The South African Constitutional Court has been instrumental in advancing Ubuntu within formal law. In *S v. Makwanyane* (1995), Justice Yvonne Mokgoro explicitly referenced Ubuntu as the philosophical foundation for dignity, compassion, and restorative justice in rejecting the death penalty. This case remains a landmark, not only for its legal implications but also for its affirmation of African values within a modern constitutional democracy (Mokgoro, 1998). It demonstrates that Ubuntu can serve as a jurisprudential guide in interpreting constitutional rights. Drawing from this precedent, national AI policies could

interpret "digital rights" not merely through the lens of autonomy and privacy, but through community well-being and social cohesion.

Embedding Ubuntu into data governance also offers possibilities for redefining core legal concepts. For instance, the principle of data ownership in most Western frameworks is framed in terms of individual consent and property rights. An Ubuntu-informed approach might shift focus toward stewardship, collective rights, and communal obligations over shared data assets. This aligns with African communitarian understandings of land, knowledge, and cultural heritage, which historically have been collectively managed. As Ifeanyi Menkiti (1984) contends, personhood in African societies is earned through fulfilling obligations to others. This logic could be applied to algorithmic systems: their legitimacy is not assumed from design principles, but must be continually earned through their impact on collective dignity.

Another strategic pathway lies in legislative preambles and constitutional values. Legal instruments could explicitly recognize Ubuntu not just as cultural heritage, but as a normative foundation for digital development. For example, the African Union's Digital Transformation Strategy (2020–2030) references inclusive development but lacks philosophical anchoring. If Ubuntu were positioned within its framework, policy interpretations could shift from efficiency and investment-readiness toward ethical, inclusive governance. This is where the African Charter on Human and Peoples' Rights also offers a tool: it already includes collective rights and duties, providing an interpretive opening for Ubuntu-aligned jurisprudence (ACHPR, 1981). This charter could be operationalized more fully in national digital strategies by insisting on obligations not just to the state, but among citizens.

On a more tactical level, national parliaments and regulatory agencies can use Ubuntu as an interpretive aid in lawmaking. Legal interpretation is not neutral; it depends on moral and political context. Just as European legal systems draw from Catholic ethics or Enlightenment rationalism, African legal systems can legitimately draw from Ubuntu. In practice, this may mean that regulators assessing AI impact evaluations would consider not just transparency reports, but also community consultations, intergenerational risks, and relational harms. Policies around biometric data, surveillance, and automated decision-making must incorporate relational accountability, a notion often absent in imported GDPR-like laws that assume atomized individuals as rights-bearers.

In regional consultations I participated in particularly in the context of drafting AI principles for Francophone Africa, it became evident that many regulators defaulted to European norms due to lack of institutional support for African alternatives. When I suggested Ubuntu as a policy principle, it was often seen as "cultural" rather than regulatory. However, through sustained dialogue and references to Mokgoro's jurisprudence, some actors began to appreciate its relevance. The discomfort was not in Ubuntu itself, but in the epistemic unfamiliarity of treating African tradition as law. This underscores the need for capacity-building among lawmakers, lawyers, and judges to develop a jurisprudence that does not separate African values from legal reasoning.

Moreover, legal pluralism in Africa already accommodates multiple normative systems, customary, religious, and statutory law co-exist. AI governance should reflect this legal ecology. A monocultural adoption of Euro-American ethics leads to what scholars call “normative misalignment” (Adeleke, 2021). Instead, governments can create AI legal codes that explicitly allow for contextual adaptation, integrating Ubuntu through public consultation and traditional authority. For example, data protection laws could include clauses that allow rural councils or traditional courts to set collective data governance norms in line with customary authority, provided they meet fundamental human rights standards.

There is also a transnational dimension. The African Continental Free Trade Agreement (AfCFTA) is already developing protocols on digital trade and data flows. Embedding Ubuntu in these frameworks would mean treating data sovereignty not just as a geopolitical matter, but as an ethical one rooted in pan-African solidarity. AI governance could be a site where Africa leads globally by showcasing a model that values shared responsibility over competition, and healing over harm.

To advance this, legal education must change. Law faculties across Africa should introduce Ubuntu jurisprudence as part of core curricula. Law students must be trained to read both H.L.A. Hart and Mogobe Ramose, both Rawls and Amadiume. Only then can Ubuntu become more than rhetoric and enter the daily legal reasoning of judges, regulators, and AI auditors.

6. The Role of the African Union and Regional Economic Communities in Ethical AI Governance

The African Union (AU) and the continent’s Regional Economic Communities (RECs) have emerged as crucial actors in shaping the legal and policy landscape for artificial intelligence. While national strategies often reflect donor priorities or domestic constraints, the AU and RECs represent sites of normative coordination and vision-setting. Their significance lies not only in their technical capacity, but in their symbolic role as custodians of African integration. To embed Ubuntu into the fabric of AI governance, these institutions must serve as vehicles for epistemic pluralism and continental sovereignty.

Historically, the AU has invoked pan-African ideals that resonate with Ubuntu’s core tenets: interdependence, collective dignity, and solidarity. From the OAU’s foundational declarations to Agenda 2063, the AU has envisioned development as both ethical and communal. In this context, the AU’s 2023 draft Continental AI Strategy is a key moment. While it references inclusion and sustainability, it remains framed largely within global discourse on rights and innovation, mirroring EU and OECD structures. This reflects what Achille Mbembe calls the “global entrapment of African futures” (Mbembe, 2016): African strategies risk merely echoing global blueprints rather than asserting endogenous norms.

Yet the AU has the potential to shift this pattern. It could formally recognize Ubuntu as a guiding ethic for digital governance, just as the European Union recognizes dignity and subsidiarity as foundational. This recognition must go beyond footnotes. It must shape how AI systems are evaluated, how risk is defined, and how policy is monitored. For example,

impact assessments under an AU model could require a communal well-being index not just algorithmic fairness scores. Licensing mechanisms could assess whether AI deployments promote relational accountability in education, agriculture, or health.

The African Commission on Human and Peoples' Rights (ACHPR), a key AU organ, already provides jurisprudential space for such shifts. Article 27 of the African Charter on Human and Peoples' Rights articulates duties owed to the community. Unlike Western rights charters, the African Charter conceptualizes personhood as fundamentally social. This foundation offers a legal scaffold for Ubuntu-based AI governance, where data practices and automated systems are measured not only by privacy protection, but by their effect on the fabric of community life (ACHPR, 1981).

RECs like ECOWAS, SADC, and the EAC are similarly positioned to champion ethical coherence. During the drafting of regional AI principles for Francophone Africa, which I participated in, it became clear that REC platforms often have closer proximity to national regulators than continental bodies do. Their value lies in facilitating regional harmonization while remaining sensitive to contextual variation. For instance, ECOWAS could establish regional AI ethics review panels drawing from traditional authorities and digital rights actors. These panels could guide cross-border deployments of AI tools in agriculture, migration control, or public health.

SADC's Digital Transformation Strategy, updated in 2022, includes provisions for data infrastructure, connectivity, and innovation ecosystems. However, its ethical components remain thin. Embedding Ubuntu would require developing regional guidance documents grounded in relational personhood. These documents should be co-authored with philosophers, traditional leaders, and civil society actors, not only technocrats or legal experts. This reflects the Ubuntu ethic of inclusive deliberation: knowledge is legitimate when it emerges from dialogic consensus, not imposed expertise (Ramose, 2003).

In practice, RECs can serve as "Ubuntu mediators" between national priorities and continental ethics. Consider how AI regulation in health might unfold. A SADC member like Zambia may wish to deploy AI triage tools. ECOWAS, having faced similar deployments, may already possess communal guidelines. The REC can broker lessons, ensure accountability, and promote regional consensus. This kind of soft law coordination builds what Claude Ake termed "development with identity" (Ake, 1996) a model rooted in cultural legitimacy and collective self-definition.

The AU and RECs also need to invest in training and epistemic infrastructure. In reviewing dozens of AI-related capacity-building programs, I found that many focus on coding, legal compliance, or funding readiness. Few teach African ethical frameworks, historical resistance movements, or oral jurisprudence. An AU-led Ubuntu Fellowship in AI governance, for example, could train policymakers to interpret data law through the lens of Bantu legal customs or Akan moral philosophy. This would not be cultural tokenism but the building of legal imagination from African starting points.

Pan-African history offers critical precedents. The Pan-African Congresses, Nkrumah's call for continental sovereignty, and the Lagos Plan of Action all centered on the notion

that development must be defined by Africans for Africans. Today's digital economy poses the same question. If Ubuntu is not embedded at the continental level, then AI governance risks becoming a new site of dependency and mimicry. The AU's AI strategy should not merely "adopt best practices." It should define them. This would align with the aspirations of the African Renaissance articulated by thinkers like Ali Mazrui and Nelson Mandela, who saw Ubuntu as a blueprint for African modernity, not a relic of pre colonial life.

Another key institution is the African Court on Human and Peoples' Rights. Though underused, it offers a judicial forum where relational harms from AI deployments could be tested. Imagine a case brought by a rural community alleging that biometric data collection for social welfare violated their communal dignity and traditional land rights. An Ubuntu-based jurisprudence would not only evaluate procedural violations but assess whether the practice disrupted societal harmony. This expands the legal terrain of digital rights.

Institutionally, however, challenges remain. Many AU and REC organs are underfunded, politicized, or driven by donor logic. To truly lead in AI governance, they must assert intellectual independence. This includes resisting "ethics washing" by external actors who seek to legitimize their technologies through superficial African partnerships. Instead, the AU should build its own ethical frameworks, invest in indigenous think tanks, and promote jurisprudential innovation rooted in lived African values.

Chapter 6: Conclusion

1. Ubuntu, Ethics, and Artificial Intelligence: Synthesis of Core Findings

This dissertation set out to investigate whether the African philosophy of Ubuntu can serve not merely as an ethical reference point but as a foundational paradigm for the governance

of artificial intelligence on the continent. From its earliest formulations, Ubuntu has offered a radically relational understanding of humanity, one that contrasts starkly with the individualist assumptions embedded in much of Western moral and political philosophy. As the preceding chapters have shown, the global AI governance framework, however well-intentioned, often universalizes values such as autonomy, privacy, and transparency without adequately interrogating the cultural and epistemic contexts from which these values emerge. By contrast, Ubuntu insists that ethical personhood is not a given but an achievement cultivated through relationships, care, and community engagement (Menkiti, 1984; Ramose, 2002).

A critical finding of this work is that African societies risk becoming passive recipients of global AI norms unless they mobilize their own philosophical resources. Through a review of existing AI frameworks (UNESCO, OECD, EU), expert interviews, and comparative ethical analysis, it has become clear that Ubuntu is not simply an alternative language, it is an alternative logic. This logic is grounded not in abstract individual rights, but in the responsibilities that arise through being embedded in social and ecological relations. As Wiredu (1996) emphasized in his work on consensus and politics, African ethics often prioritize communal harmony and dialogical reasoning, principles that offer powerful correctives to the adversarialism and utilitarian calculus dominating current AI ethics regimes.

Throughout this dissertation, reflections from fieldwork, summits, expert dialogues, and academic literature were woven together to illustrate how Ubuntu can be both a moral compass and a policy tool. For instance, in interviews with experts such as Dr. Melody Musoni and Benjamin Akinmoyeje, as well as insights from Kweya Chechi's remarks at the FieldFisher podcast during Africa Week, it became evident that ethical AI in Africa must engage with questions of digital inclusion, intergenerational justice, and linguistic diversity. The Microsoft Africa team's use of AI for climate resilience, agricultural diagnostics, and language preservation affirms that AI's promise on the continent cannot be fulfilled by importing technical solutions alone, it requires ethical imagination rooted in African lived realities.

Moreover, reflections shared in the African AI summits in Kampala, Rwanda, and Paris reinforced the view that the continent's approach to AI must be both globally connected and locally situated. In Lusaka, discussions during the Digital Rights and Inclusion Forum (DRIF) raised critical concerns about surveillance, sovereignty, and data militarization. In Riyadh at the Global IGF, the African delegation stressed the need for frameworks that are not only interoperable but also epistemically inclusive. These insights informed this research's stance that Ubuntu is not a supplement to Western ethics but a paradigm that recasts the debate entirely.

Importantly, Ubuntu's relevance is not confined to normative ideals. Its implications for institutional design, legal reasoning, and public policy are substantial. By framing ethics as a matter of interdependence, Ubuntu challenges the current siloed approaches to AI governance. It demands participatory structures where elders, youth, engineers, regulators, and traditional custodians co-create AI policies. This is especially pertinent in a continent marked by diverse epistemologies, oral traditions, and plural legal systems. As Mogobe

Ramose (1999) has argued, Ubuntu is a jurisprudence as much as it is a philosophy, its logic permeates how justice is understood, how conflict is resolved, and how legitimacy is claimed.

The research has also shown the dangers of epistemic injustice when African traditions are relegated to footnotes in policy debates. In chapter two, we examined how the concept of autonomy as it appears in Kantian ethics, and by extension in AI policy documents, often erases relational understandings of agency that are more prevalent in African thought. As Fricker (2007) has argued, the failure to recognize alternative knowledge systems constitutes a form of epistemic violence. This is not merely symbolic; it influences how technologies are designed, deployed, and governed.

Against this backdrop, the integration of Ubuntu into AI governance offers an emancipatory horizon. It allows African policymakers to reclaim the ethical space, to assert that African philosophies are not peripheral but central to discussions about the future of human-machine interaction. This also responds to the critiques of Afro-pessimism and technological determinism. As highlighted in previous reflections, this thesis is animated by a sense of Afro-optimism, an insistence that Africa is not merely a testbed for imported innovation but a laboratory of alternative futures.

Still, the journey is not without tension. Ubuntu resists codification. It is not a checklist or compliance protocol. It is a dynamic, lived ethic that must be continuously interpreted. This makes its translation into policy and law both challenging and essential. As articulated in meetings with Professor Souleymane Bachir Diagne, Ubuntu's power lies not in being a fixed doctrine but in its capacity to animate new forms of solidarity, care, and ecological responsibility.

Thus, the research concludes that Ubuntu is not only suitable for AI ethics, it is necessary. It provides African governments with a normative foundation that is both rooted in heritage and responsive to the digital present. It enables ethical governance that is not imported or imposed but cultivated through dialogue, reflection, and shared humanity. In a world increasingly governed by algorithms, Ubuntu reminds us that ethics begins not with the self, but with the other.

2. Policy and Institutional Recommendations for Africa

Building on the findings of this research, concrete recommendations are now needed to translate Ubuntu's normative power into functional AI governance strategies. These recommendations are directed primarily at African governments, intergovernmental bodies, regulatory agencies, and research institutions. Their aim is to offer pathways for institutionalizing ethical governance frameworks that are rooted in the continent's own philosophical traditions, rather than externally imposed norms.

First, there is an urgent need to establish national and regional AI ethics commissions that integrate traditional and indigenous knowledge holders. These commissions should not only include computer scientists, data analysts, and legal experts, but also philosophers, cultural historians, elders, and community leaders who are versed in local ethical reasoning.

This recommendation is grounded in reflections from the AI-in-Africa Landscape Study (GCG, 2023), which underscores the risks of narrow technocratic policy making. In AI governance, excluding epistemic diversity is not merely a failure of representation, it is a structural flaw that leads to misaligned policy. In interviews with Melody Musoni and in exchanges during the Africa AI Summit, the value of incorporating Afrocentric ethics into AI advisory bodies was repeatedly emphasized.

Second, regulatory innovation must include experimentation with Ubuntu-inspired AI sandboxes. These are flexible, context-specific legal environments where AI technologies can be tested in accordance with community-derived ethical norms. Unlike conventional regulatory sandboxes focused solely on market innovation, Ubuntu sandboxes would prioritize social cohesion, communal rights, and intergenerational equity. This idea emerged through discussions during the ICANN82 meetings in Seattle, where it became evident that Western-style regulation often fails to respond to the realities of digital development in Africa's hybrid legal environments. Ubuntu-based sandboxes would permit digital experimentation while maintaining ethical oversight grounded in customary law and participatory governance.

Third, AI ethics education must be embedded in African academic curricula from secondary school through postgraduate levels. This involves a double move: integrating courses on African philosophical traditions and critically engaging with dominant Western AI frameworks. As Nkiru Nzegwu (2006) and Ifi Amadiume (1987) have argued, education in Africa must serve as a site of intellectual decolonization. In the context of AI, this means teaching young Africans not only how to code, but also how to think ethically from within their own traditions.

Fourth, digital policy must be harmonized across the continent through the active involvement of the African Union and Regional Economic Communities (RECs), but in a way that avoids flattening cultural distinctions. While the AU's Continental AI Strategy marks an important step, national implementation must allow space for linguistic, religious, and philosophical diversity. The ECOWAS Digital Strategy 2020-2023 and the Southern African Development Community (SADC) Cybersecurity Model Law offer promising precedents. However, as noted in the Microsoft Africa Week podcast, standardization must not lead to homogenization. Instead, regional bodies should serve as platforms for pluriversal negotiation, spaces where varied African moral vocabularies can find expression.

Fifth, national and regional public consultation processes must go beyond symbolic inclusion. Too often, so-called multistakeholder consultations are dominated by donor priorities, with limited community engagement. Drawing on personal experience organizing AI policy consultations in Lomé, it became clear that civil society actors, youth leaders, and traditional institutions are frequently sidelined. One way to address this is to establish national digital ethics assemblies, modeled on traditional forums of deliberation. These could convene citizens, elders, policymakers, and technologists to co-develop guidelines for AI deployment in sectors such as health, agriculture, justice, and education.

Sixth, the notion of “consent” in AI deployment must be reframed through the lens of relational accountability. Current AI ethics standards, including those promoted by the OECD and UNESCO, often reduce ethical safeguards to individual data consent. This fails to capture the communal nature of harm in African societies. As discussed in interviews and field reflections, AI-driven projects in rural settings often impact entire clans or land-based communities, even when consent is obtained from a single individual. Ethical AI governance in Africa must adopt collective impact assessments, rooted in customary norms about land, ancestry, and spiritual stewardship. This requires regulators to go beyond the European GDPR model and co-develop Afrocentric data protection frameworks.

Seventh, investment in ethical AI must be matched by investment in infrastructure and local innovation. Ethical governance cannot be sustained if African countries remain dependent on foreign cloud providers, proprietary platforms, and outsourced technical expertise. The recommendations of the Smart Africa Alliance (2022) emphasize the importance of data sovereignty and infrastructure localization. Drawing from reflections by Strive Masiyiwa during our interview, Africa’s AI governance must be accompanied by tangible support for local developers, open-source platforms, and community-led tech solutions.

Finally, the role of the diaspora must be embraced not as a symbolic bridge, but as a substantive force in ethical AI innovation. Diasporic African thinkers, engineers, and ethicists many of whom are already active in global institutions should be formally invited to contribute to national AI roadmaps. This recommendation draws from patterns observed in Ghana’s inclusion of its diaspora in strategic planning, and echoes calls made during the Global AI Summit in Rwanda for Afro-descendant inclusion in global tech governance. Ubuntu ethics extends beyond geographical Africa. It is a moral vision that connects all those who see humanity as relational, dignity as shared, and future-making as collective.

In sum, Ubuntu offers not only a moral philosophy but a governance architecture. It provides African states with a vocabulary for ethical deliberation that is neither derivative nor exclusionary. These recommendations offer a blueprint for translating Ubuntu into concrete policy, regulatory, and educational instruments. As technology continues to reshape society, Africa’s ethical frameworks must not remain reactive. They must become constitutive of how AI is imagined, built, and governed from the ground up.

3. The Global Relevance of Ubuntu-Inspired AI Governance

As the ethical governance of artificial intelligence becomes a matter of global urgency, the significance of African thought systems, particularly Ubuntu, is no longer limited to regional frameworks. The proposition advanced throughout this dissertation is not that Ubuntu should be “added” to existing ethical models, but rather that it offers an alternative paradigm for AI governance rooted in relationality, moral interdependence, and shared responsibility. This concluding section examines the potential influence of Ubuntu on global AI discourse and policy making, positioning Africa not as a recipient of ethical norms, but as a co-creator of global standards.

Contemporary AI ethics is dominated by Euro-American philosophical assumptions that valorize autonomy, individual rights, and procedural fairness. As argued in earlier chapters, these concepts, while not without merit, fail to account for the embedded, collective, and historical nature of moral decision-making in many parts of the world. Ubuntu intervenes by offering a framework in which ethics begins not with the isolated subject, but with the social fabric that binds persons together. This resonates with calls by scholars such as Boddington (2017) and Floridi (2013) for more holistic approaches to AI ethics that consider the broader consequences of AI technologies on social cohesion and community well-being.

One area where Ubuntu may have global relevance is in redefining the notion of harm. In most AI ethical frameworks, harm is assessed in relation to the individual, whether data is misused, whether autonomy is violated, or whether discrimination has occurred. Ubuntu, however, encourages a relational view of harm. An action is unethical not only because it causes individual injury but because it fractures communal harmony. This understanding could shift the basis of risk assessments in AI systems, prompting developers and regulators to evaluate technologies based on their effects on relationships, social cohesion, and collective trust. For instance, biometric surveillance may be deemed lawful and non-invasive in Western contexts if consent is given. But from an Ubuntu perspective, even consensual surveillance that engenders fear or undermines social trust could be viewed as morally corrosive.

The global debate on AI and fairness also stands to benefit from Ubuntu-informed perspectives. Mainstream approaches often define fairness in terms of statistical parity or equal treatment, using mathematical tools to measure bias. Yet these approaches can miss deeper injustices rooted in historical exclusion, unequal starting points, or the loss of dignity. Ubuntu fairness is contextual and dialogical. It prioritizes restorative action, mutual recognition, and the repair of broken relations. This could inspire international institutions, including those within the United Nations system, to rethink fairness beyond technical metrics and adopt processes of deliberative justice that center voices from marginalized communities.

Moreover, Ubuntu ethics has implications for the governance of AI beyond national or continental levels. It suggests a model of multilateralism based not on geopolitical competition but on solidarity, reciprocity, and shared responsibility. These values align with emerging calls for a Digital Commons, shared infrastructures and principles to govern digital technologies globally. In conversations with Professor Souleymane Bachir Diagne, he emphasized that Ubuntu is not a parochial ethic, but a cosmopolitan one. It calls for “making humanity together,” a vision that disrupts the binary between the local and the global. As countries negotiate instruments such as the Global Digital Compact and frameworks under WSIS+20, African voices informed by Ubuntu can offer models of collective stewardship that resonate far beyond the continent.

The challenges to such global influence are, however, significant. Many African frameworks are still seen as derivative or normative supplements, rather than epistemic sources. The dominance of English-language discourse, Western funding logics, and policy gatekeeping mean that Ubuntu may be included symbolically but excluded substantively.

This necessitates not only intellectual advocacy but institutional reform, changing how think tanks, international organizations, and governance platforms recognize and validate ethical models from the Global South.

Despite these obstacles, the tide is shifting. The UNESCO Recommendation on the Ethics of AI (2021) makes notable references to cultural pluralism and the importance of local knowledge systems. The African Union's Continental Strategy on AI includes Ubuntu explicitly as a value. Civil society organizations such as Research ICT Africa and the Mozilla Foundation have also begun funding projects that integrate indigenous ethics into AI policy. These trends indicate a growing appetite for ethical pluralism and a readiness to move beyond tokenistic inclusion.

For global institutions, the value of Ubuntu lies not only in normative guidance but in methodological insight. Ubuntu ethics is not codified in rigid rules. It is lived, narrated, and performed through relationships. Applying Ubuntu in AI governance thus requires a shift in how ethical deliberation itself is structured, from abstract rule-making to dialogical engagement; from technical audits to community-based reflection. This approach could inform new methodologies for technology impact assessments, user-centered design, and participatory policymaking, particularly in areas like health data, education algorithms, or environmental AI.

In closing, the Ubuntu philosophy represents more than an ethical supplement to dominant paradigms. It is a foundational challenge to the logic of AI governance as it stands today. It critiques the fragmentation of ethical subjects, the abstraction of social life, and the reduction of values to metrics. It calls instead for a situated, communal, and intergenerational ethic that sees technology as embedded in the moral architecture of society. This dissertation has aimed to show that African traditions, when treated not as artifacts but as philosophical interlocutors, can reconfigure global debates about how we live with intelligent machines.

If taken seriously, Ubuntu-inspired AI governance could influence everything from regulatory principles to global standards-setting. It could democratize not only access to technology but access to the ethical narratives that shape our digital future. Africa, long viewed as a latecomer to technological revolutions, may thus become a moral frontrunner offering the world not only tools but also wisdom for governing them.

References

Acemoglu, D., & Johnson, R. (2023). Rebalancing artificial intelligence. *Finance & Development*, International Monetary Fund.

Adeleke, F. (2021). Legal pluralism and digital governance in Africa. *African Journal of International and Comparative Law*, 29(3), 412–430.

African Commission on Human and Peoples' Rights. (1981). *African Charter on Human and Peoples' Rights*. Banjul.

African Union. (2020). *Digital Transformation Strategy for Africa (2020–2030)*. Addis Ababa: African Union Commission.

African Union. (2022). *Data Policy Framework*. Addis Ababa: African Union Commission.

African Union. (2024). *Continental Strategy for Artificial Intelligence*. Addis Ababa: African Union Commission.

African Union Development Agency NEPAD. (2024). *AI policy and innovation ecosystems in Africa*. Midrand.

AI4D Africa. (2021). *Responsible artificial intelligence for development in Africa*. Ottawa: International Development Research Centre.

AI4D Africa. (2022). *Community-centred AI governance models*. Nairobi.

Akinmoyeje, B. (2025). Interview with Emmanuel Agbenonwossi on AI ethics and policy sovereignty in Africa. Unpublished transcript.

Alcoff, L. (2006). *Visible identities: Race, gender, and the self*. Oxford: Oxford University Press.

Amadiume, I. (1987). *Male daughters, female husbands*. London: Zed Books.

Arakpogun, E. O., Elsahn, Z., Olan, F., & Elsahn, F. (2021). Artificial intelligence in Africa: Challenges and opportunities. In *The Fourth Industrial Revolution*. Cham: Springer.

Autor, D., Levy, F., & Murnane, R. (2003). The skill content of recent technological change. *Quarterly Journal of Economics*, 118(4), 1279–1333.

Azeroual, F. (2024). *Artificial intelligence in Africa: Challenges and opportunities*. Rabat: Ministry of Digital Transition and Administration Reform.

Berg, A., Papageorgiou, C., & Vaziri, F. (2023). The uneven effects of technological change. *Finance & Development*, International Monetary Fund.

Birhane, A. (2021). Algorithmic injustice. *Patterns*, 2(2).

Birhane, A., & Guest, O. (2020). Towards decolonising computational sciences. *Patterns*, 1(8).

Boddington, P. (2017). *Towards a code of ethics for artificial intelligence*. Cham: Springer.

- Brynjolfsson, E., & Unger, G. (2023). The macroeconomic implications of artificial intelligence. *Finance & Development*, International Monetary Fund.
- Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of the Conference on Fairness, Accountability and Transparency*, 77–91.
- Cabral, A. (1973). *Return to the source*. New York: Monthly Review Press.
- Cath, C., Wachter, S., Mittelstadt, B., Taddeo, M., & Floridi, L. (2018). Artificial intelligence and the “good society”. *AI & Society*, 33(4), 505–515.
- Chechi, K. (2025). Responsible AI and inclusive digital governance in Africa. Interview on *FieldFisher Africa Week Podcast*. Unpublished audio.
- Coetzee, P. H., & Roux, A. P. J. (2002). *The African philosophy reader*. London: Routledge.
- Cornell, D., & Van Marle, K. (2005). *Exploring Ubuntu*. Oxford: Hart Publishing.
- Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonizing human life*. Stanford: Stanford University Press.
- Crawford, K. (2021). *Atlas of AI*. New Haven: Yale University Press.
- Creemers, R., Triolo, P., & Webster, G. (2022). China’s algorithmic governance. *New America*.
- de Sousa Santos, B. (2014). *Epistemologies of the South*. Boulder: Paradigm Publishers.
- Diagne, S. B. (2018). *En quête d’Afrique(s)*. Paris: Albin Michel.
- Diagne, S. B. (2021). *De langue à langue: L’hospitalité de la traduction*. Paris: Albin Michel.
- Diagne, S. B. (2025). Interview with Emmanuel Agbenonwossi on Ubuntu and ethical pluralism. Unpublished transcript.
- Diop, C. A. (1981). *Civilisation ou barbarie*. Paris: Présence Africaine.
- Eubanks, V. (2018). *Automating inequality*. New York: St. Martin’s Press.
- European Commission. (2021). *Proposal for a Regulation on Artificial Intelligence (AI Act)*. Brussels.
- European Commission. (2024). *Artificial Intelligence Act*. Brussels.
- Fanon, F. (1963). *The wretched of the earth*. New York: Grove Press.

- Floridi, L. (2013). *The ethics of information*. Oxford: Oxford University Press.
- Floridi, L., & Cows, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1).
- Fricker, M. (2007). *Epistemic injustice*. Oxford: Oxford University Press.
- Garcia, J. M., & Kelly, T. (2020). *Artificial intelligence for Africa*. Washington DC: World Bank.
- Gillwald, A., Mothobi, O., & Rademan, B. (2018). The state of ICT in Africa. *Research ICT Africa*.
- Gobodo-Madikizela, P. (2020). *A human being died that night*. Boston: Houghton Mifflin.
- Green, B. (2021). The contestation of tech ethics. *Big Data & Society*, 8(2).
- Gwagwa, A., Kazim, E., et al. (2020). Responsible AI in Africa. *Philosophy & Technology*, 33, 563–576.
- Hagendorff, T. (2020). The ethics of AI ethics. *Minds and Machines*, 30, 99–120.
- Harding, S. (2004). *The feminist standpoint theory reader*. New York: Routledge.
- Hountondji, P. J. (1997). *Endogenous knowledge*. Dakar: CODESRIA.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389–399.
- Kant, I. (1997). *Groundwork of the metaphysics of morals* (original work published 1785). Cambridge: Cambridge University Press.
- Kurbalija, J. (2025). Meeting notes on AI governance and ethics. Unpublished internal memo.
- Landemore, H. (2023). Inclusive democracy and AI. *Finance & Development*, International Monetary Fund.
- Letseka, M. (2012). In defence of Ubuntu. *Studies in Philosophy and Education*, 31(1), 47–60.
- Mandela, N. (1994). *Long walk to freedom*. Boston: Little, Brown.
- Masakhane. (2022). *Community-led NLP for African languages*. Masakhane Research Reports.
- Masiyiwa, S. (2025). Interview with Emmanuel Agbenonwossi on infrastructure and AI sovereignty. Unpublished transcript.

- Metz, T. (2011). Ubuntu as a moral theory. *African Human Rights Law Journal*, 11, 532–559.
- Metz, T. (2017). *African values, human rights and group rights*. Oxford: Oxford University Press.
- Mignolo, W. (2009). *The darker side of Western modernity*. Durham: Duke University Press.
- Milan, S., & Treré, E. (2019). Big data from the South. *Television & New Media*, 20(4), 319–335.
- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1, 501–507.
- Mokgoro, Y. (1998). Ubuntu and the law in South Africa. *Potchefstroom Electronic Law Journal*, 1(1).
- Ngũgĩ wa Thiong’o. (1986). *Decolonising the mind*. London: James Currey.
- Nzegwu, N. (2006). *Family matters*. Albany: SUNY Press.
- OECD. (2019). *Principles on Artificial Intelligence*. Paris.
- Ramose, M. B. (2002). *African philosophy through Ubuntu*. Harare: Mond Books.
- Taylor, L. (2017). What is data justice? *Big Data & Society*, 4(2).
- Tuck, E., & Yang, K. W. (2012). Decolonization is not a metaphor. *Decolonization*, 1(1).
- Tutu, D. (1999). *No future without forgiveness*. New York: Doubleday.
- UNESCO. (2021). *Recommendation on the Ethics of Artificial Intelligence*. Paris.
- UNESCO. (2023). *AI readiness assessment in Africa*. Paris.
- Veale, M., & Borgesius, F. Z. (2021). Demystifying the EU AI Act. *Computer Law Review International*, 22(4), 97–112.
- Wiredu, K. (1996). *Cultural universals and particulars*. Bloomington: Indiana University Press.
- World Bank. (2023). *The size and distribution of digital connectivity gaps in Sub-Saharan Africa*. Washington DC.
- Wynter, S. (2003). Unsettling the coloniality of being. *CR: The New Centennial Review*, 3(3), 257–337.

Appendices

Appendix A

Semi-Structured Interview Guide

Thematic block	Purpose	Sample guiding questions
Background and role	Situate expertise and professional engagement	What is your current role in AI governance or digital policy? What ethical issues arise most frequently in your work?
Global AI ethics frameworks	Assess perceptions of dominant frameworks	How do you assess UNESCO, OECD, or EU AI governance frameworks? Which assumptions travel poorly into African contexts?
African ethical traditions	Explore Ubuntu and communal ethics	How do you understand Ubuntu as an ethical framework for AI governance? What does it change in how harm and responsibility are defined?
Policy and institutional challenges	Identify governance and capacity constraints	What are the main institutional barriers to ethical AI governance in Africa?
Legal dimensions	Explore feasibility of legal integration	Can relational or communal ethics be embedded into law and regulation?
Continental and regional governance	Examine AU and REC roles	What role should the African Union and Regional Economic Communities play in AI governance?
Forward-looking perspectives	Capture future priorities	What would an Ubuntu-informed AI governance model look like in practice over the next 5–10 years?

Appendix B

Expert Interviews Conducted for the Study

Interview code	Role / expertise	Affiliation type	Geographic scope	Format
INT-01	Philosopher, AI ethics	University-based	West Africa / Global	Online
INT-02	Senior policy advisor	Continental institution	Continental	In person
INT-03	Data governance specialist	Regional body	Southern Africa	Online
INT-04	Legal scholar	University-based	East Africa	Online
INT-05	Digital rights advocate	Civil society organisation	West Africa	In person
INT-06	AI researcher	Research institute	Southern Africa	Online
INT-07	Technology policy specialist	International organisation	Global / Africa	Online
INT-08	Innovation ecosystem lead	National agency	East Africa	In person
INT-09	Private sector executive	Technology firm	Pan-African	Online
INT-10	Gender and technology specialist	NGO / research network	West Africa	Online
INT-11	Data protection advisor	National authority	Southern Africa	Online
INT-12	AI and development researcher	Think tank	East Africa	Online

INT-13	Public sector strategist	Government ministry	West Africa	In person
INT-14	Continental digital policy expert	AU-affiliated programme	Continental	Online
INT-15	Ethics and governance scholar	University-based	Global / Africa	Online
INT-16	Infrastructure specialist	Telecom / infrastructure group	Pan-African	Online
INT-17	Youth and innovation advocate	Innovation hub	East Africa	In person
INT-18	AI regulation consultant	Independent expert	Southern Africa	Online
INT-19	Human rights lawyer	Legal advocacy organisation	West Africa	Online
INT-20	Digital inclusion specialist	International NGO	Africa	Online
INT-21	Government digital strategist	National government	East Africa	In person
INT-22	AI ethics researcher	Research centre	Southern Africa	Online
INT-23	Diaspora policy expert	Policy network	Global / Africa	Online
INT-24	Technology economist	Multilateral organisation	Global	Online
INT-25	Civil society coalition coordinator	Regional digital rights coalition	West Africa	Online
INT-26	AI governance thought leader	Independent / academic	Pan-African	Online

Appendix C

Policy and Governance Documents Reviewed

Level	Institution / actor	Document	Year
Global	UNESCO	Recommendation on the Ethics of Artificial Intelligence	2021
Global	OECD	OECD Principles on Artificial Intelligence	2019
Global	European Union	Regulation (EU) 2024/1689 – Artificial Intelligence Act	2024
Global	United Nations	Global Digital Compact	2024
Global	IEEE	Ethically Aligned Design (First Edition)	2019
Global	World Economic Forum	The AI Governance Journey: Development and Opportunities	2021
Continental	African Union	Digital Transformation Strategy for Africa (2020–2030)	2020
Continental	African Union	Data Policy Framework	2022
Continental	African Union	Convention on Cyber Security and Personal Data Protection	2014
Continental	African Commission on Human and Peoples' Rights	African Charter on Human and Peoples' Rights	1981
Regional	ECOWAS	Supplementary Act A/SA.1/01/10 on Personal Data Protection	2010
Regional	Smart Africa Alliance	Digital Identity Blueprint	2020
Regional	Smart Africa Alliance	Digital Health Blueprint	2024
Multistakeholder	Partnership on AI	Annual Report	2022