
AI LITERACY FELLOWSHIP FOR AFRICAN POLICYMAKERS · COHORT 1

Eight Capstones, One Country

AI Policy Projects for the Republic of Ubuntia

The Republic of Ubuntia is a composite country invented for this fellowship, so that fellows from across the continent could confront one national problem together without singling out any real state.

EDITOR-IN-CHIEF

Najeeb G. Abdulhamid

CONTRIBUTING EDITORS

Abdulrahman Hayatu

Abba Muhammad Gadanya

Abubakar Manga

Alhassan Haruna Umar

Salim Ibrahim

Sani Isah Abba

Bashir I. Dodo

Fatima Abdulkadir Jidda

Ibrahim Labaran Ali

Khalil Yusuf Adamu

Abubakar Salisu Abdullahi

Eight Capstones, One Country

AI Policy Projects for the Republic of Ubertia

PUBLISHED BY

OpenSchool Initiative

Abuja, Nigeria · First edition, 2026

EDITOR-IN-CHIEF

Najeeb G. Abdulhamid

CONTRIBUTING
EDITORS

Abdulahman Hayatu	Fatima Abdulkadir Jidda
Alhassan Haruna Umar	Abubakar Salisu Abdullahi
Bashir I. Dodo	Abubakar Manga
Khalil Yusuf Adamu	Sani Isah Abba
Abba Muhammad Gadanya	Ibrahim Labaran Ali
Salim Ibrahim	

THEMATIC GROUP
MENTORS

Named with their affiliations in the Editorial Board and Mentors section overleaf.

AUTHORS

Forty-one fellows across eight working groups, named in the Contributors section.

BASED ON

The eight capstone projects submitted by Cohort 1 of the AI Literacy Fellowship for African Policymakers, October 2025 to February 2026.

RIGHTS

© 2026 OpenSchool Initiative
<https://doi.org/10.5281/zenodo.22066741>

STATUS

An open educational and policy-learning resource. The Republic of Ubertia and its institutions, data and scenarios are fictional. This volume is not adopted policy, nor legal, clinical or financial advice. Anyone adapting it must verify local law, mandates, evidence and costs.

CC BY 4.0

FREE TO COPY,
ADAPT AND SHARE

Licensed under Creative Commons Attribution 4.0 International. You may copy, adapt and share this work for any purpose, provided you credit the original and indicate any changes. Third-party material remains subject to its own rights. The DOI above is a concept DOI: it always resolves to the latest version.

CITE THE VOLUME AS

Abdulhamid, N. G., Hayatu, A., Umar, A. H., Dodo, B. I., Adamu, K. Y., Gadanya, A. M., Ibrahim, S., Jidda, F. A., Abdullahi, A. S., Manga, A., Abba, S. I., & Ali, I. L. (Eds.). (2026). *Eight Capstones, One Country: AI Policy Projects for the Republic of Ubertia*. OpenSchool Initiative. <https://doi.org/10.5281/zenodo.22066741>

Contents

FRONT MATTER

Credits and rights.....	ii
Editorial board and mentors.....	iv
About this volume.....	vi
A note from the Editor-in-Chief.....	x
Contributors.....	xiii

THE CAPSTONE PROJECTS

01 Cross-Cutting Policy and Digital Sovereignty.....	1
<i>National AI Governance and Digital Sovereignty Blueprint</i>	
02 Education and Skills for the Future.....	18
<i>National AI in Education Policy</i>	
03 Labour, Employment and Future of Work.....	35
<i>AI and the Future of Work Policy Blueprint</i>	
04 Finance, Trade and Inclusion.....	55
<i>Responsible Fintech and Digital Inclusion Framework</i>	
05 Agriculture and Food Systems.....	72
<i>AI Policy Blueprint for Climate-Smart Agriculture and Food Security</i>	
06 Climate Resilience and Urban Systems.....	90
<i>AI Policy Blueprint for Climate-Resilient, Data-Driven Urban Systems</i>	
07 Health and Wellbeing.....	113
<i>Responsible AI for Health Equity</i>	
08 Justice, Rights and Governance.....	135
<i>A Justice, Rights and Governance AI Framework</i>	

Each chapter follows the nine-section blueprint set out in the Capstone Project Handbook and closes with its references, a verification of its claims, and a record of the editorial judgements applied to it.

Editorial board

Najeeb G. Abdulhamid · Editor-in-Chief

OpenSchool Initiative.

Abba Muhammad Gadanya

MEng Electrical Engineering; COREN certified; member, NSE and NATE; certified by NEMSA and RENAC (Germany). Power systems, project management and community outreach.

Abdulrahman Hayatu

BSc Computer Science (SLUK); Google Certified Educator. Cybersecurity and machine learning.

Alhassan Haruna Umar

PhD Electrical Engineering (UTM); MSc Electrical/Electronic Engineering (University of Leeds); BEng Electrical and Computer Engineering (FUT Minna); Microsoft Innovative Educator and Google Workspace Administrator. Physiological sensing and processing.

Bashir I. Dodo

PhD Computer Science (Brunel University London); MSc Computer Systems Engineering (UEL); BSc Software Engineering; Associate Fellow, HEA. Medical image analysis, policy and online education.

Khalil Yusuf Adamu

BEng Electrical/Electronic Engineering; MEd Educational Administration and Planning; MBA Management; COREN Technologist; OSHA certified safety supervisor. Energy systems, education and training, partnerships and fundraising.

Abubakar Manga

BEng Electrical/Electronic Engineering; co-founder and former COO, FlexiSAF Edusoft. Business strategy and institutional development, education technology operations, digital transformation.

Salim Ibrahim

Doctoral candidate; MBA Innovation, Startups and Entrepreneurship (VIGNAN'S, Andhra Pradesh, India); MCL (Bayero University, Kano); BSc Ed (ABU, Zaria); NCE; certified quality assurance assessor and financial literacy trainer.

Fatima Abdulkadir Jidda

Systems IT support specialist; MSc and PGD Information Technology; HND Computer Science; member, Computer Professionals Registration Council of Nigeria (CPN).

Abubakar Salisu Abdullahi

MSc Software Engineering (Nile University of Nigeria); Bachelor of Information Technology (IUEA); IADNCS, IDITC (Informatics Academy). Software security, binary analysis and intelligent vulnerability detection.

Sani Isah Abba

Assistant Professor of Civil Engineering, Prince Mohammad Bin Fahd University, Saudi Arabia. Water and environmental engineering, climate change, and machine learning algorithms.

Ibrahim Labaran Ali

Lecturer, The University of Warwick. Artificial intelligence, data analytics and digital transformation. Also served as a thematic group mentor.

Thematic group mentors

Each working group was supported by a mentor who reviewed its problem definition, evidence and policy design as the capstone developed.

Ibrahim Labaran Ali

Lecturer, The University of Warwick. Artificial intelligence, data analytics and digital transformation.

Njeri Ngaruiya Ng'ang'a

ICT4D researcher, educator and consultant.

Gilles Quentin Hacheme

Senior Research Scientist, Microsoft AI for Good Lab.

Shamsuddeen Hassan Muhammad

Google DeepMind Academic Fellow; Advanced Research Fellow, Imperial College London.

Andrew Abu

Technology entrepreneur; Founder, Young Innovators of Nigeria (YIN).

Fauziya Ado Yakasai

Certified data scientist and academic. Statistics, machine learning and applied mathematics.

Felermينو Dário Mário António Ali

Postdoctoral Fellow, Microsoft Research Africa; doctoral candidate in Computer Engineering, University of Porto, Portugal.

Aliyu Ahmad

Senior Azure Cloud Engineer, Data and AI, Microsoft.

Editorial note. Names and qualifications are published as supplied by the Editor-in-Chief, with spelling and punctuation normalised for consistency. No credential has been altered. Members of the editorial board are listed here in that capacity; the offices they hold within OpenSchool Initiative are not given, because they are not the basis on which they appear. Honorifics and professional titles are likewise not attached to contributor names anywhere in this volume: the roster mixes academic doctorates, professional registration and contributors who use no title, and no single scheme covers all three. Each person's qualifications are stated in full beneath their name instead.

About this volume

This volume collects the eight capstone projects produced by the first cohort of the AI Literacy Fellowship for African Policymakers. Each was written by a team of fellows working on a different sector of the same country — a country that does not exist. What follows explains who convened the fellowship, what the capstone asked of its teams, and why an invented republic turned out to be the most honest setting available for the work.

OpenSchool Initiative

OpenSchool Initiative is a volunteer-powered, Africa-focused non-profit that helps individuals and institutions build practical digital and AI capability for learning, work and public value. It began during the COVID-19 lockdown as a response to interrupted education, offering free digital learning that supported remote teaching, digital literacy, work readiness and career preparation. It has since grown into a body of inclusive, evidence-led classes, workshops and learning communities serving learners and institutions across the continent.

Its work joins information and media literacy, responsible AI literacy, AI for work and productivity, teacher development, community digital inclusion and evidence-informed learning design. The fellowship sits inside that wider mission. It was built to move AI literacy past awareness and tool use, toward the institutional capacity that responsible adoption actually requires.

The AI Literacy Fellowship for African Policymakers

The fellowship was a four-month, fully virtual, pan-African programme running from 20 October 2025 to 13 February 2026. It brought together people whose work touches public policy: government, academia, research, civil society, journalism, professional practice and institutional leadership.

Its purpose was not to explain artificial intelligence. It was to build the practical, evidence-based capacity to interrogate AI claims, weigh risk against opportunity, design policy responses, communicate trade-offs honestly, and govern adoption under African conditions.

Fellows moved through four modules: AI foundations and Africa's position in the emerging landscape; ethics, risk and governance; adoption, innovation and capacity building; and AI for development, closing with capstone delivery. Lectures, workshops, masterclasses, self-paced study, mentoring and peer collaboration were combined so that the programme would end in applied policy work rather than in conceptual understanding.

The capstone project

The capstone was the fellowship’s applied policy laboratory. Fellows were organised into eight thematic teams, each taking a different sector of national life: cross-cutting policy and digital sovereignty; education and skills; labour and the future of work; finance, trade and inclusion; agriculture and food systems; climate resilience and urban systems; health and wellbeing; and justice, rights and governance.

Each team moved from problem definition and evidence review through policy design, institutional responsibilities, implementation choices, safeguards and indicators, to a finished sector blueprint. Mentors reviewed the work as it developed, and peers challenged it. The eight projects in this volume are those completed submissions.

Why Ubuntia

The fellowship’s concept note of May 2025 set out the ambition for actionable capstone work. When the learning experience was designed, the organisers ran into a practical and ethical problem: how could fellows drawn from many African countries and professional settings work together on a single national policy problem without making one real country the default?

Choosing a real country would have privileged one national experience and made every other context peripheral. It would also have invited classroom analysis to be read as criticism of, or advocacy for, a particular government.

That risk was not abstract. Some fellows work inside public institutions, advise them, or report on them. Frank examination of weak capacity, contested law, institutional failure, political trade-offs or unsafe deployment could easily be mistaken for an official position, or for a judgement about a fellow’s own country. A shared fictional setting created room for serious and uncomfortable analysis without exposing fellows, their employers or their governments to political attribution.

Ubuntia was created for that purpose: a hypothetical composite African republic serving as a common policy canvas. It gathers plausible demographic, economic, social, technological, governance and environmental conditions found across the continent, but it does not represent, rename or conceal any single real state. Every team received the same baseline, which made it possible to argue about hard choices, test competing approaches and propose reform without implying that any capstone spoke for a real government.

The shared canvas did something else as well. Because all eight teams worked on different sectors of the same country, their proposals could be set against one another, their institutional overlaps resolved, and the whole tested for coherence. Eight unrelated papers would not have permitted that.

How the shared canvas worked

1. Begin with one authoritative country profile carrying the shared demographic, institutional, economic and social baseline.
2. Assign each thematic team a distinct sector challenge within that same national setting.
3. Allow teams to introduce further assumptions or proposals, while distinguishing those additions from the original baseline.
4. Bring the sector outputs together, resolve institutional overlaps, and test whether they form a coherent national architecture.
5. Treat real African countries as later adaptation environments, where every recommendation must be retested against local conditions.

BOUNDARY

Ubuntia is not a disguised version of any country, a substitute for local evidence, or a policy ready for adoption. It is a teaching and policy-simulation environment, designed to support inclusive collaboration, guard against false political attribution, and make cross-sector synthesis possible. Anyone adapting this work to a real jurisdiction must test every recommendation against local law, evidence, capacity, culture and public participation.

The Republic of Ubuntia at a glance

These parameters were given to every team, and recur across all eight chapters. Where a chapter introduces a figure of its own, the editors have said so in that chapter's verification table.

Population	42 million
Median age	20 years
Rural population	Over 60%
Internet penetration	About 48% nationally, lower in rural districts
Youth unemployment	Over 30%
Agriculture	About 60% of the workforce
Capital	Lumora
Technology hub	Kifaru City
Data protection	Data Protection Act, 2021

How this volume is organised

Each of the eight chapters follows the structure set out in the Capstone Project Handbook: executive summary; context and rationale; problem analysis; policy objectives; proposed interventions; governance and implementation model; risk and mitigation plan; monitoring and evaluation framework; and conclusion. Submissions that arrived in a different order have been reorganised to it, so that a reader moving between sectors finds the same argument in the same place.

Every chapter closes with the same editorial apparatus: its references, an editorial note on sources, a verification of the claims it makes, and a table of the editorial judgements applied to it. That last table is unusual in a published volume, and deliberate. It records what the editors changed and why, including the places where a defect was resolved on the evidence of the chapter's own text rather than left as a query.

A companion volume

A separate publication, the National AI Governance Blueprint for the Republic of Ubutia, draws on these eight capstones to construct a single national governance architecture, reconciling the sectors into one institutional design. This volume is the source material: the eight sector projects as their teams wrote them, edited for consistency and accuracy but not merged. The two are intended to be read together.

Source basis. AI Literacy Fellowship for African Policymakers Concept Note, May 2025; the Cohort 1 replication playbook; the National AI Governance Blueprint for the Republic of Ubutia; the Republic of Ubutia case study; and the eight submitted capstone projects.

A note from the Editor-in-Chief

This volume presents eight sector policy projects as a work of record. The editorial task was to bring them to a common standard of structure, evidence and presentation without displacing the judgement of the teams who wrote them.

Preparing capstone policy work for publication raises a question that does not arise in ordinary academic editing. These chapters were written as applied policy design inside a taught programme, and they are published after that programme has closed. There is therefore no author review cycle: no opportunity to return a query and receive an answer. An editor in that position must decide, in advance and in the open, what standard of intervention is legitimate.

The standard applied

Three principles governed the edit.

Structure before style. Each chapter was set to the nine-section blueprint prescribed by the Capstone Project Handbook, from executive summary through to conclusion, so that a reader moving between sectors meets the same argument in the same place. Where a submission was organised differently, its material was resequenced rather than rewritten.

Evidence must resolve. Every locator in this volume was opened and confirmed to lead to the work cited. A reference that cannot be followed is not a reference, and in a permanently indexed record it is a liability.

Interventions must be visible. Any editorial decision a reader might reasonably contest is recorded in the chapter where it was made, alongside what the submission originally said.

The scope of intervention

Within those principles, editorial work fell into six categories.

Structure. Resequencing to the Handbook's nine sections, with cross-references and numbering brought into line.

Text integrity. Restoration of characters, words and headings lost in file conversion, always against the surrounding text. Where an original reading could be recovered it was recovered; no wording was substituted for convenience.

Tables. Reassembly of tables that conversion had dissolved into loose paragraphs, and their return to the prescribed six-column intervention format.

Completeness. Where the Handbook structure required a component a chapter did not carry — a risk and mitigation plan, in one case — it was constructed from that chapter’s own stated risks, guiding principles and interventions, and declared in the text as constructed.

References. Sixty-two references appear across the volume, the majority located by the editors for chapters that carried no reference list, with in-text citations added at the points where each chapter relies on them.

Figures. Twenty-six figures appear here. Artwork carrying errors set into the image itself, or supplied below print resolution, was redrawn in the volume’s house style with its content unchanged.

The boundary

The rule that governed every one of those categories was this: consolidate, never invent. Material could be moved, rejoined, renumbered, repaired or reformatted. It could not be created.

No policy position was altered. No target, figure or recommendation was invented, softened or strengthened. Where a chapter’s argument is weak, it remains weak and visible; the editorial function is to present an argument accurately, not to improve it. Where a table had a cell the authors did not fill, it reads “not specified by authors” rather than carrying a plausible substitute.

The same restraint governed attribution. Where an affiliation was not supplied, the entry says so, because the editors had no basis on which to infer an institution and would not guess at one.

Transparency as method

Every chapter closes with a table headed Editorial judgements, setting out where the editors intervened, what the submission said, and what was done. There are seventy-nine such entries across the eight chapters.

A record of this kind is uncommon in a published volume. It is included because an edited text that conceals its edits asks the reader to take the editor on trust, while a text that records them invites scrutiny. Any reader who disagrees with a judgement can see precisely what was changed and reconstruct what was submitted.

Each chapter also carries a verification table separating claims internal to the Ubuntia case study from those resting on external evidence, so that scenario parameters are never mistaken for findings.

The discipline of restraint

Editorial confidence is not the same as editorial correctness, and the difference is worth illustrating.

One chapter cites an International Labour Organization compendium and dates it 2026. The date appeared wrong; the process it described ran through the previous year, and the obvious editorial move was to correct it. Checking the source showed the authors were right — the compendium was issued in January 2026, following a presidency that concluded in 2025. The citation stands as written, and the episode is recorded in that chapter’s judgement table in the authors’ favour.

Verification cut in the other direction as well. Sources that looked entirely credible were discarded when they failed to resolve: a dead DOI, a withdrawn page, and one report that proved not to exist at all. Where a publisher’s server refused automated access, the record was confirmed through the Crossref metadata registry rather than assumed.

WHERE DOUBT REMAINS VISIBLE

One chapter states that over sixty per cent of digital public projects fail to scale. The figure could not be traced to a primary source. It has not been removed, because it may well be sound; and it has not been given a citation, because none could be verified. It stands marked as not verified. That is the position this volume takes wherever certainty was not available.

A closing word

The policy thinking in these chapters is coherent, evidenced and institutionally literate, and the editorial work recorded here was overwhelmingly a matter of structure, verification and presentation. The arguments are the authors’ own.

They belong to forty-one fellows across thirty-four institutions and fourteen countries, supported by eight mentors, and they are published here to be read, argued with and adapted.

Najeeb G. Abdulhamid

Editor-in-Chief · OpenSchool Initiative, Abuja

Contributors

Cohort 1 of the AI Literacy Fellowship for African Policymakers graduated fifty-five fellows. Forty-one are named in this volume.

Authorship here follows substantial contribution to a submitted capstone project, which is a narrower test than completing the fellowship. Each fellow is credited as an author of the chapter their team produced, on the basis of the names recorded on that submission. Graduating from the fellowship and authoring a published chapter are distinct achievements, and the absence of a name from this list says nothing about the first.

Contributors are listed by thematic area, in the order the chapters appear, in eight thematic teams, with the affiliations they gave during the fellowship. The same names appear at the opening of each chapter.

01 Cross-Cutting Policy and Digital Sovereignty

National AI Governance and Digital Sovereignty Blueprint

Luvuyo Ngcobo · University of Pretoria, South Africa

Abdul Mohammed Isa · Voice of Nigeria, Nigeria

Buhari Umar Ali · Independent National Electoral Commission, Nigeria

Fadumo Ladan Suleiman · Somalia Stream, Somalia

Kassim Umar · Voice of Nigeria, Nigeria

Uloma Okoro · University of East London, Nigeria

Abubakar Sadiq Yushau · Jigawa State Polytechnic, Dutse, Nigeria

02 Education and Skills for the Future

National AI in Education Policy

Abderrahim Chalfaouat · Hassan II University of Casablanca, Morocco

Abdiaziz Bashiir · DeepSom Institute of AI and Technology, Somaliland

Ahmed Esmat · The American University in Cairo, Egypt

Genet Abebe Gabre · Kotebe University of Education, Ethiopia

Hajer Ben Hadj Salem · High Institute of Human Sciences of Tunis, Tunisia

Nasiru Aliyu · Centre for Geodesy and Geodynamics, National Space Research and Development Agency, Toro, Nigeria

Naziru Mikail Abubakar · University of Westminster, London, United Kingdom

Nouruddeen Bashir Umar · Aliko Dangote University of Science and Technology, Wudil, Nigeria

Said Ali Said Yunus · State University of Zanzibar, Tanzania

03 Labour, Employment and Future of Work

AI and the Future of Work Policy Blueprint

Shuaibu Sambo Othman · Voice of Nigeria, Nigeria

Hadiza N. Tijjani · Nigerian Television Authority, Nigeria

Lorna Okatch · Government of Kenya, Kenya

Sarah Ibrahim · Cairo University, Egypt

Aisha Usman Muhammad · Gombe State University, Nigeria

04 Finance, Trade and Inclusion

Responsible Fintech and Digital Inclusion Framework

Rushmi Christine Matete · Jackfruit Associates Limited, Kenya

Fridah Jebowen Keitany · Chamber Business Academy, Kenya National Chamber of Commerce and Industry, Kenya

Sani Bako Mohammed · Central Bank of Nigeria, Nigeria

05 Agriculture and Food Systems

AI Policy Blueprint for Climate-Smart Agriculture and Food Security

Abdelhafeez Elemam Gasoum · Affiliation not supplied

Abubakar Hassan Inuwa · Foundation for Sustainable Smallholder Solutions (FSSS), Nigeria

Hassan Faruk Maradun · Usmanu Danfodiyo University Sokoto, Nigeria

Idris Muhammad Shamsudeen · Central Bank of Nigeria, Nigeria

Suleiman Usman Yusuf · Affiliation not supplied

Tukur Garba Arab · Voice of Nigeria, Nigeria

06 Climate Resilience and Urban Systems

AI Policy Blueprint for Climate-Resilient, Data-Driven Urban Systems

Ismail Musa Umar · Nigerian Television Authority, Nigeria

Maryam Rabi Aliyu · Cyprus International University, Mersin, Türkiye

07 Health and Wellbeing

Responsible AI for Health Equity

Samuel Sanaa Brobbey · Ministry of Health, Ghana

Ijeoma Chiemela · VitalNutrient Nutrition Foundation, Nigeria

Genet Shanko Dekebo · Adama Science and Technology University, Ethiopia

Edna C. Too · Chuka University, Kenya

Zubair Usman · Jigawa State Ministry of Health, Dutse, Nigeria

08 Justice, Rights and Governance

A Justice, Rights and Governance AI Framework

Abubakar Usman Abubakar · Universiti Malaya, Malaysia

Muhammad Nasambo · Nigeria Revenue Service, Nigeria

Oluwaseyi Opeyemi Jimoh · Osun State Ministry of Innovation, Science, Technology and Digital Economy, Nigeria

Sadiya Ibrahim Galadima · National Judicial Institute, Nigeria

A note on this list. Contributor names are given without honorifics or professional titles throughout this volume; the roster spans academic doctorates, professional registrations and contributors who use no title, and no single scheme covers all three. Personal email addresses supplied with the submissions have been withheld: publishing them in a permanently indexed record would expose contributors to unsolicited contact for the life of the deposit. Institutional affiliation is retained in every case where it was supplied.

Cross-Cutting Policy and Digital Sovereignty

National AI Governance and Digital Sovereignty Blueprint

Establishing a National AI Commission and Public Algorithm Registry Within 36 Months

CASE STUDY

The Republic of Ubuntuia

WORKING GROUP

Cross-Cutting Policy and Digital Sovereignty Working Group (Group 8)

Luvuyo Ngcobo · University of Pretoria, South Africa

Abdul Mohammed Isa · Voice of Nigeria, Nigeria

Buhari Umar Ali · Independent National Electoral Commission, Nigeria

Fadumo Ladan Suleiman · Somalia Stream, Somalia

Kassim Umar · Voice of Nigeria, Nigeria

Uloma Okoro · University of East London, Nigeria

Abubakar Sadiq Yushau · Jigawa State Polytechnic, Dutse, Nigeria

Suggested citation: Ngcobo, L., Isa, A. M., Ali, B. U., Suleiman, F. L., Umar, K., Okoro, U., & Yushau, A. S. (2026). National AI Governance and Digital Sovereignty Blueprint. In N. G. Abdulhamid (Ed.), *Eight Capstones, One Country* (pp. xx–xx). OpenSchool Initiative. <https://doi.org/10.5281/zenodo.22066741>

1. Executive Summary

The Republic of Ubuntuia is a youthful Sub-Saharan democracy of 42 million citizens, and the country currently stands at a transformative juncture. With 60% of its population living in rural areas and a median age of twenty, Ubuntuia's demographic advantage presents a powerful opportunity for inclusive innovation. However, this promise is threatened by deep structural imbalances in connectivity, institutional capacity, and digital dependency.

The current AI landscape in Ubuntuia is marked by fragmented pilot projects, donor-driven systems poorly adapted to local languages and contexts, and rising concerns about sovereignty, exclusion, and trust. Without a national strategy, Ubuntuia risks becoming a testing ground for AI systems built for foreign contexts, potentially worsening inequality, reinforcing bias, and deepening reliance on non-African platforms.

For this reason, this National AI Governance and Digital Sovereignty Blueprint presents a homegrown, human-centred approach to managing AI in the public interest. It is grounded in principles of fairness, participation, and sovereign infrastructure, and proposes a focused objective: to establish a National AI Commission and a Public Algorithm Registry within a 36-month timeframe. These two institutions form the governance backbone for inclusive, ethical, and transparent AI deployment in Ubuntuia. Its value proposition is to confront institutional fragmentation and public mistrust, and to invest in legal, ethical, and participatory AI governance. Through this, Ubuntuia can create a coherent framework that ensures AI systems serve national priorities and empower all citizens.

1.1 Six Strategic Interventions

1. Legal establishment of the National AI Commission with oversight powers
2. A national AI Ethics Framework and mandated impact assessments for high-risk tools
3. Capacity-building for public officials, youth, and marginalized communities
4. Sectoral AI Focal Units to integrate AI in health, education, agriculture, and justice
5. A transparent, publicly accessible Algorithm Registry and grievance mechanisms
6. Broad-based citizen engagement in AI design and oversight

These measures will bridge the gap between policy and practice, reduce reliance on external systems, and position Ubuntuia to lead ethically in Africa's digital transformation. The interventions align with continental frameworks such as the African Union's Digital Transformation Strategy (2020–2030) and regional calls for algorithmic accountability and digital sovereignty. If implemented with urgency, this blueprint will provide a scalable model of responsible innovation rooted in equity, trust, and democratic values.

2. Context and Rationale

Ubuntia’s development journey presents both challenges and opportunities. With a population of 42 million and a median age of twenty, the country stands among the youngest nations in Sub-Saharan Africa. This demographic dividend is paired with a growing technology ecosystem centred in Kifaru City, and emerging regional partnerships that position Ubuntia to lead on digital innovation. However, institutional fragility, uneven development, and deepening technological dependency threaten to undermine this potential. This blueprint therefore stems from an urgent need to harness AI as a national development instrument, while safeguarding sovereignty, promoting inclusion, and building public trust in technology.

2.1 Ubuntia’s Development Landscape

Over 60% of Ubuntians live in rural areas where agriculture dominates livelihoods. Although primary school enrollment is high, dropout rates remain significant, especially for girls, and access to tertiary education is limited. Urbanization is accelerating, with Lumora and Kifaru City absorbing thousands of young migrants each year, many of whom face high youth unemployment and underemployment. Startups are flourishing in urban centres, often backed by diaspora capital, but rural communities remain digitally excluded and underserved by infrastructure and services.

Ubuntia’s digital and economic landscape also has its share of challenges. As mobile money and fintech companies expand rapidly, cyber fraud and weak consumer protection measures have become major areas of concern. Digital literacy is uneven, and local governments remain under-resourced, which affects service delivery and data collection. AI pilot projects — especially in agriculture, health, and finance — have been mostly donor-funded and imported, lacking adaptation to local languages such as Ubuntu, and failing to align with cultural norms and national priorities. These systems often perform poorly in real-world conditions and may reinforce inequality by excluding women, rural citizens, and non-English or non-French speakers.

At the governance level, Ubuntia passed a Data Protection Act in 2021, signalling a shift toward rights-based digital governance. However, enforcement remains weak and regulatory capacity is thin. The lack of a national AI strategy or regulatory framework leaves the country vulnerable to unregulated AI deployments, many of which operate without transparency or ethical oversight. As recent analyses of African AI environments note, countries without coordinated policies risk “becoming experimental sites for foreign innovation” that neither reflects nor respects local needs.

2.2 Global AI Transformations and the African Context

AI is rapidly transforming global economies, redefining sectors from agriculture to healthcare to public administration. Countries that have adopted AI strategies early, such as Rwanda, Kenya, and Nigeria, are already seeing gains in government efficiency, innovation ecosystems, and regional competitiveness. However, global AI development has often sidelined African priorities. A significant share of AI models are trained on datasets from the Global North, leading to performance failures and ethical blind spots when deployed in African contexts. Without deliberate policy action, Ubuntia risks importing these harms. The African Union’s Digital Transformation Strategy (2020–2030) emphasizes the need for digital sovereignty, ethical governance, and regional harmonization to ensure AI supports inclusive development.

Building on this, the blueprint offers a vision rooted in homegrown AI governance, aligned with regional principles while tailored to Ubuntu's unique socio-economic and linguistic realities.

2.3 Rationale for a Human-Centred, Sovereignty-Driven Approach

AI cannot be treated as a neutral technology; it reflects the values, assumptions, and power structures embedded in its design and deployment. In fragile governance environments like Ubuntu, where institutional trust remains low, public oversight is limited, and policy enforcement is uneven, the unchecked adoption of AI systems risks deepening existing inequalities. A human-centred, sovereignty-driven approach is therefore a practical necessity. Without deliberate safeguards, algorithmic systems may replicate existing social hierarchies, marginalize underrepresented communities, and concentrate decision-making power in opaque and unaccountable systems.

In Ubuntu's context, this risk is already visible. Farmers in the Central Highlands rely on informal messaging networks for weather alerts due to the absence of reliable public platforms; rural teachers still depend on chalkboards while ed-tech pilots remain confined to urban centres; fintech users in urban markets face rising cyber fraud amid regulatory gaps. These realities are structural signals of exclusion. Without AI systems designed for these conditions, national deployment risks reinforcing a digital divide where innovation is reserved for the few and governance bypasses the many.

The blueprint therefore adopts data sovereignty as a foundational pillar. Evidence from regional policy studies shows that national control over data infrastructure, local algorithmic standards, and context-specific datasets is essential for safeguarding self-determination in the digital age. Without this control, Ubuntu risks falling into "algorithmic dependency" — a form of digital colonization in which foreign actors design, deploy, and govern technologies that affect citizens without accountability or cultural relevance. This dependency is especially dangerous when imported systems trained on foreign datasets fail to accommodate local languages, legal frameworks, or ethical norms, as has already been observed in donor-led AI pilots in agriculture and health.

Moreover, the cost of inaction is immediate and cumulative. If Ubuntu does not establish a coordinated national AI governance framework, the following risks become increasingly probable:

- **Widening inequality**, as AI tools entrench biases based on gender, geography, and socio-economic status
- **Opaque governance**, where public sector decisions are delegated to unregulated algorithms
- **Sovereignty threats**, as critical infrastructure becomes reliant on foreign-controlled systems
- **Stalled innovation**, as local ecosystems struggle to scale without regulatory clarity or public trust

For these reasons, the blueprint advocates urgent, collaborative action. A human-centred and sovereignty-driven AI governance framework will be institutionally strategic, socially responsive, and necessary to ensure that AI contributes to inclusive development rather than reinforcing historical inequities.

3. Problem Analysis

Ubuntia's AI ecosystem is currently fragmented, reactive, and highly vulnerable to external influence. In the absence of a national strategy or coordinated regulatory framework, AI deployments are largely driven by foreign donors and multinational vendors, resulting in systems poorly adapted to local languages, infrastructure, and social norms. An internal review of 25 pilot AI tools across agriculture, health, and education found that over 80% were not trained on Ubuntian data or evaluated for local impact. This raises serious concerns about effectiveness and equity.

The institutional architecture to manage these risks is also weak. Although Ubuntia passed a Data Protection Act in 2021, enforcement remains minimal due to limited regulatory staffing and capacity. Fewer than 30% of ministries have designated data officers, and inter-ministerial platforms for data sharing remain underdeveloped. This leads to siloed systems, especially in education, health, and agriculture, where AI uptake is increasing without corresponding oversight. There are no formal grievance mechanisms for citizens harmed by algorithmic decisions, further weakening public trust in digital services.

The structural digital divide compounds these challenges. While mobile penetration in urban areas exceeds 85%, rural connectivity lags below 40%, limiting equitable access to digital tools and platforms. Public sector data shows that only 12% of rural schools have any form of internet access, and less than 8% of national AI training initiatives reach women or girls. These disparities are mirrored in the technology workforce, where women represent fewer than 20% of AI-related roles.

Four Structural Challenges to AI Governance in Ubuntia




Figure 1.1 Four structural challenges to AI governance in Ubuntia.

In this vacuum, Ubuntia risks becoming a passive consumer of foreign-developed AI systems — tools that may entrench inequality, automate public functions without accountability, and expose sensitive infrastructure to external control. Without urgent policy intervention, these gaps will widen, weakening the country's capacity to govern its digital future and eroding public trust in the very systems intended to support development.

4. Policy Objectives

This policy aims to establish a legally mandated National AI Commission and a Public Algorithm Registry as the central pillars of Ubutia's AI governance framework. The Commission will provide cross-sectoral coordination, ethical oversight, and regulatory clarity, while the Registry will ensure transparency, accountability, and public engagement in all government-deployed AI systems.

Table 1.1 Institutional milestones over the 36-month implementation period.

Institution	Milestone	Deadline
National AI Commission	Formally established, staffed, and funded	Month 12
Public Algorithm Registry	Launched with all government AI tools registered and a functioning grievance mechanism	Month 30
Both institutions	Fully operational	Month 36

The rationale for prioritizing this objective is as follows:

- The Commission will serve as a central structure for coordinating AI policy and deployment across sectors.
- It will ensure that all public AI tools meet standards of safety, fairness, and alignment with national values.
- The Public Algorithm Registry will give citizens visibility into how AI is used in public services and provide channels for redress.

Together, these institutions will form the backbone of a sovereign, trusted, and inclusive AI ecosystem. To avoid stifling innovation, the Commission will adopt a tiered audit system that prioritizes high-risk tools while allowing low-risk AI innovations, especially from startups, to proceed with fast-track review protocols and innovation sandboxes by Year 2.




Figure 1.2 Framework for national AI governance, with the Commission and Registry at its centre.

5. Proposed Interventions

To achieve the objective of establishing a legally mandated National AI Commission and Public Algorithm Registry, Ubuntia will pursue strategic interventions across five mutually reinforcing domains. These actions will ensure that, within 36 months, the country has the institutional, legal, ethical, technical, and public engagement infrastructure necessary to regulate AI in the public interest. The interventions are sequenced to allow phased implementation: laying a strong foundation in the first year, scaling operations in the second, and consolidating systems and public participation by the third.

5.1 Legal and Institutional Set-Up

The first priority is the legal establishment of the National AI Commission (NAIC) through an enabling Act of Parliament. This Commission will serve as the apex institution for AI regulation, coordination, and oversight. The legal framework will define its mandate, powers, composition, and relationship with sectoral agencies.

Alongside legislative efforts, a Technical Secretariat will be established within the Ministry of Digital Innovation and Data Sovereignty to manage day-to-day implementation. This Secretariat will lead the operationalization of the NAIC, including the development of staffing structures, funding mechanisms, and internal protocols. This intervention is foundational and will be completed within the first 12 months, creating the institutional backbone for all subsequent activities.

To ensure long-term sustainability beyond the initial 36 months, the Commission will be embedded as a permanent statutory body with a dedicated budget line in the national appropriation bill. It will also explore supplementary revenue through licensing fees for high-risk AI deployments and certification of AI providers.

5.2 Ethical Oversight and Standards

To ensure that AI tools deployed in Ubuntia are aligned with national values, the NAIC will develop a national AI Ethics Framework in consultation with ethicists, technologists, civil society, and vulnerable communities. The framework will provide principles and criteria for fairness, accountability, cultural sensitivity, and harm prevention.

In parallel, all AI systems used in public services will be subject to Algorithmic Risk and Impact Assessments before deployment. These assessments will examine bias, safety, privacy, and societal impact. They will be mandated by law and administered by the NAIC's Ethics Council. Together, these instruments will operationalize ethical standards and ensure that all public AI is subject to meaningful scrutiny and pre-emptive risk mitigation. Initial guidelines will be issued by Month 18, with mandatory compliance required by Year 3.

Similar governance models have been successfully adopted elsewhere. Rwanda's National AI Policy emphasizes ethical auditing through a central authority, while Kenya's digital transformation strategy integrates data protection and AI regulation under a unified framework. The European Union's AI Act includes a tiered risk-based registry system that informs this blueprint's approach. These examples demonstrate the practicality and effectiveness of establishing central oversight with context-specific safeguards.

5.3 Capacity Building and Education

To enable effective governance and citizen oversight, Ubuntuia will invest in AI capacity across three groups: public servants, youth and women, and the general public. A National AI Literacy Programme will deliver foundational AI knowledge through schools, radio broadcasts, mobile platforms, and community hubs, with special attention to rural areas and content in the Ubuntu language. Civil servants involved in digital service delivery will undergo compulsory training in ethical AI procurement, data management, and regulatory compliance. These trainings will begin in Year 1 and scale nationally by Year 2.

To complement training efforts, a national AI Innovation Grant Scheme will be launched to fund local startups, researchers, and community-based technology labs. This will promote the development of context-relevant AI solutions by Ubuntians, for Ubuntuia. Over time, it will strengthen digital sovereignty by reducing reliance on imported technologies.

5.4 Sectoral Integration

The success of the NAIC depends on its ability to coordinate with existing ministries. AI Focal Units will be established in priority sectors — health, education, agriculture, and justice — to translate national standards into sector-specific practice. These units will identify relevant AI use cases, co-design AI tools with communities, and support departments in complying with ethical standards and registry requirements. In the Ministry of Health, for instance, the focal unit will guide the use of AI for maternal health screening while ensuring compliance with national ethics guidelines and participation frameworks. By Year 2, these sectoral units will be active and feeding into the NAIC’s reporting and oversight systems, operationalizing the Commission’s mandate across government.

In the agricultural sector, for example, AI can be applied to improve early warning systems for crop failure using satellite data and weather patterns. The first tangible “win” will be the rollout of a government-certified AI tool that provides SMS-based pest outbreak alerts to maize farmers in the Central Highlands by the end of Year 2.

5.5 Citizen Engagement

To build trust and legitimacy, citizens must be able to understand, question, and influence AI systems that affect their lives. The Public Algorithm Registry will be a digital platform listing all government-deployed AI systems, their purposes, data sources, and contact points for public feedback. Ubuntuia will also establish grievance and redress mechanisms linked to the registry, allowing individuals to report harm, bias, or misuse of AI tools. These reports will feed directly into the NAIC’s monitoring and auditing cycles.

Open consultations and community co-design sessions will be hosted across regions to involve citizens in shaping AI systems. By Year 3, every AI tool used in public services must be both listed in the Registry and subject to citizen feedback and ethics review.

The Public Algorithm Registry will be maintained as core national infrastructure, with long-term technical support contracted to local firms and regular updates managed by the Commission’s Digital Operations Unit, ensuring its continuity beyond the project phase.

6. Governance and Implementation Model

To successfully deliver Ubuntu's National AI Commission and Public Algorithm Registry within the 36-month period, a clear and inclusive governance structure will be adopted. This model balances national oversight with local responsiveness and promotes collaboration across all sectors of society.




Figure 1.3 Governance and implementation model, centred on the National AI Commission.

6.1 National AI Commission

The NAIC will be the lead institution overseeing AI regulation and coordination in Ubuntu. It will be established by law and operate as an independent body, reporting annually to Parliament. Its core functions will include:

- Approving and registering all public sector AI tools
- Overseeing ethical standards and algorithmic risk assessments
- Coordinating with ministries and partners
- Publishing regular public reports

6.2 Technical Implementation Secretariat

Housed within the Ministry of Digital Innovation and Data Sovereignty, the Secretariat will serve as the operational arm of the NAIC. It will:

- Lead the day-to-day rollout of programmes
- Support ministries with AI deployment
- Manage training, procurement, and public outreach
- Coordinate timelines and budgets

6.3 AI Focal Units in Key Ministries

Dedicated AI units will be created in major sectors such as health, education, agriculture, and justice. These units will:

- Identify practical AI use cases
- Ensure compliance with national ethics and transparency standards
- Collect data and report back to the Secretariat and the NAIC
- Work with communities to co-design relevant tools

6.4 AI Ethics Council

An expert advisory group under the NAIC will guide ethical oversight. It will include ethicists, legal experts, community leaders, and technologists. Its role will be to:

- Review sensitive or high-risk AI projects
- Host public consultations
- Recommend updates to ethical guidelines

6.5 Citizen Participation Channels

The public will have multiple ways to engage with the AI system:

- A Public Algorithm Registry providing open access to all AI tools used by government
- A grievance platform allowing citizens to report concerns
- Regular town halls and digital consultations held across the country

This model ensures that AI governance in Ubuntuia is transparent, accountable, and inclusive — from lawmaking to local communities. It provides a realistic and people-centred pathway to manage AI in the public interest.

7. Risk and Mitigation Plan

Table 1.2 *Risks, potential impacts and mitigation strategies.*

#	Risk	Potential Impact	Mitigation Strategy
1	Institutional resistance	Delays in implementation and weak coordination	Early stakeholder engagement and mandatory training for ministries
2	Legislative delays	Late establishment of the AI Commission and Registry	Work with parliamentary champions and prepare interim executive mechanisms
3	Low public trust	Poor citizen uptake and limited feedback	National awareness campaigns and co-designed digital tools
4	Privacy and data risks	Breach of sensitive data or misuse of registry information	Strong cybersecurity and clear data protection protocols
5	Technical skill gaps	Weak deployment, monitoring, and ethics compliance	Training for public servants and embedded technical support in ministries
6	Funding challenges	Slow rollout or reduced scope of interventions	Diversify funding sources and align programmes with donor and diaspora support
7	Vendor dependence	Loss of sovereignty or misaligned systems	Enforce ethical procurement rules and prioritize local developers
8	Regional misalignment	Reduced participation in digital trade or data-sharing platforms	Align national actions with African Union and AfCFTA frameworks

8. Monitoring and Evaluation Framework

To ensure accountability, learning, and timely delivery of the AI governance objective, Ubuntia will adopt a results-based monitoring and evaluation framework. This system will track progress across all implementation phases, identify challenges early, and promote transparency.

8.1 Guiding Principles

- **Results-oriented.** Focused on outputs and outcomes, not just activities.
- **Participatory.** Inclusive of civil society, youth, women, and rural communities.
- **Transparent.** Open dashboards and public reporting to build trust.
- **Adaptable.** Flexible enough to adjust plans based on feedback and evidence.

8.2 Tools and Mechanisms

- **Quarterly dashboards.** Each implementing agency will track key metrics such as AI literacy enrollment, ethics reviews completed, and registry usage.
- **Annual AI governance report.** Led by the National AI Commission and presented to Parliament, using disaggregated data.
- **Public feedback channels.** Citizens can submit experiences with AI tools via the registry and community forums.
- **Ethics audits.** The AI Ethics Council will conduct yearly reviews of selected AI deployments.

8.3 Sample Indicators

Table 1.3 Sample indicators by objective area.

Objective Area	Key Indicators
Legal and institutional setup	AI Commission operational; law passed within Year 1
Ethics and oversight	Percentage of public AI tools with completed impact assessments
Capacity building	Number of public officers and youth trained; percentage of female participants
Public engagement	Number of AI tools listed in the registry; number of grievance cases resolved
Sector integration	Number of functional AI focal units across ministries

8.4 Evaluation Timeline

Table 1.4 Evaluation timeline.

Period	Evaluation activity
Year 1	Baseline assessments, dashboard development, first ethics audit
Year 2	Mid-term review, public report, course correction
Year 3	Final evaluation, impact report, and sustainability roadmap

9. Conclusion and Call to Action

Ubutia's digital trajectory will be defined by the decisions it makes today. This blueprint articulates a clear, achievable path — anchored in the establishment of a National AI Commission and a Public Algorithm Registry that transforms fragmented governance into a coherent system of oversight, ethics, and trust. This is not merely an institutional reform; it is a sovereignty imperative. By investing in coordinated regulation, participatory mechanisms, and national standards, Ubutia affirms its right to shape AI on its own terms, aligned with its values, languages, and developmental priorities.

We call on Parliament to adopt the enabling legislation with urgency, and on ministries to embed this framework across sectors. Civil society and the private sector must serve not only as watchdogs but as co-creators of ethical AI. Regional and international partners are urged to align with Ubutia's agenda for sovereign, inclusive innovation.

The window for passive adaptation is closing. With deliberate leadership and cross-sectoral commitment, Ubutia can define not only its digital future, but influence the continent's path toward equitable AI governance.

References

This chapter was submitted with its own reference list and numbered in-text citations. The editors have normalised it to APA 7th style, removed three duplicated entries, repaired URLs that had been broken by stray spaces, and verified every locator. Where the authors supplied no DOI, one has been added.

- African Union Commission. (2020). *The digital transformation strategy for Africa (2020–2030)*. African Union. <https://au.int/en/documents/20200518/digital-transformation-strategy-africa-2020-2030>
- Agbeno, E. K. (2025). Algorithmic sovereignty: Regulating AI and data colonialism under the AfCFTA framework. *International Journal of Human Movement, Education and Legal Studies*. <https://ijhmeels.com/index.php/journal/article/view/7>
- Ajuzieogu, U. C. (2025). *The African Union AI continental strategy and the development–governance paradox*.
- Gwagwa, A., Kraemer-Mbula, E., Rizk, N., Rutenberg, I., & De Beer, J. (2020). Artificial intelligence (AI) deployments in Africa: Benefits, challenges and policy dimensions. *The African Journal of Information and Communication*, 26, 1–28. <https://doi.org/10.23962/10539/30361>
- Mensah, G. B. (2025). *Digital sovereignty in the age of AI: Reconciling national data control with cross-border AI innovation*. ResearchGate. <https://www.researchgate.net/publication/387959552>
- Salami, A. O. (2024). Artificial intelligence, digital colonialism, and the implications for Africa’s future development. *Data & Policy*, 6, e67. <https://doi.org/10.1017/dap.2024.62>
- Savage, S., & Savage, L. (2025). *Inclusive and secure AI: A global perspective on policy and technical developments*.
- Yusuf, B. (2025). *Robust digital governance frameworks in Africa*. Centre for International Governance Innovation.
- Zimba, M., Jouini, M., & Ndaka, A. K. (2025). Context-aware Africa-led designing of responsible artificial intelligence technologies. In *Trustworthy AI: African perspectives* (pp. 145–164). Springer Nature Switzerland.

Editorial note on sources

This chapter carries the volume’s most developed scholarly apparatus, with numbered in-text citations throughout, and the editors have preserved every source the authors chose. The submitted list contained eleven numbered entries, of which three were duplicates: Salami (2024) appeared as both reference 1 and reference 6, Ajuzieogu (2025) as both 3 and 7, and the African Union Commission strategy as both 9 and a further unnumbered entry. These have been consolidated into single entries and the list set alphabetically, as APA requires.

Three URLs had been corrupted by stray spaces inserted through the document — the African Union link, for example, was submitted as “[https:// au.int/en /d ocu ments/2 0200 210 /d igital - tran sfor mati on - strategy- africa](https://au.int/en/d ocu ments/2 0200 210 /d igital - tran sfor mati on - strategy- africa)”. All have been repaired and checked. DOIs have been added for the Salami and Gwagwa articles, which the submitted list gave without locators. Two entries — Ajuzieogu (2025) and Savage and Savage (2025) — were submitted without a publisher, journal or locator, and no published version could be located; they are recorded as the authors gave them.

Verification of claims

Table 1.5 Editorial verification of claims.

Claim as written	Section	Editorial verification
“Over 80% of 25 reviewed pilot AI tools were not trained on Ubuntian data”	§3	Internal to the case study, presented as the finding of an internal review within the scenario.
“Fewer than 30% of ministries have designated data officers”	§3	Internal to the case study. No external source is cited.
“Rural connectivity below 40%” and “urban mobile penetration exceeds 85%”	§3	Internal to the case study. Note that other chapters in this volume give national internet penetration as 48%, which is consistent with a split of this kind between urban and rural areas.
“Only 12% of rural schools have internet access”; “fewer than 20% of AI-related roles held by women”	§3	Internal to the case study. No external source is cited.
“The Data Protection Act of 2021”	§2.1, §3	Internal to the case study. The same instrument is referenced in Thematic Areas 2, 3 and 7, indicating a shared scenario parameter rather than a real statute.
Rwanda, Kenya and the European Union as comparative governance models	§5.2	Externally supported. Rwanda’s National AI Policy, Kenya’s digital transformation strategy and the EU AI Act all exist and operate broadly as described.
Digital colonialism and algorithmic dependency	§2.3	Externally supported by the sources the authors themselves cite — Salami (2024) and Agbeno (2025) — both verified by the editors.

Editorial judgements

No author review cycle was available before publication. Where the submitted material contained defects whose resolution could be established with confidence from the surrounding text, the volume editors have resolved them rather than leaving them in the published record. Each such resolution is listed below. No substantive claim, target, or policy position has been altered.

Table 1.6 Editorial judgements applied to this chapter.

Location	As submitted	Editorial resolution
Figure 1.1 (submitted map)	The national map labelled three separate cities “Kwanza”, repeated “Weak Data Regulation” across two different zones, and marked Kifaru City as the capital with a capital-city symbol.	The capital marking contradicts the shared case study, in which Lumora is the capital and Kifaru City the technology hub — as this chapter itself states in §2.1. Rather than reproduce these errors, the four challenge categories the map conveyed have been redrawn as a diagram, populated with the figures the authors give in §3. Geographic context is provided by the map in Thematic Area 3.
Figure 1.3 (submitted governance model)	The artwork rendered “Housed in the Ministry” as “Host irin the Ministry” and “Health, Education, Agriculture, Justice” as “...Justice for me”. Resolution was also too low for print.	Redrawn at print resolution with the text corrected from the chapter’s own §6.2 and §6.3.
Figure 1.2 (submitted framework)	Content was accurate but resolution was too low for print.	Redrawn at print resolution. Content unchanged.
§1, Executive Summary	A stray “>” character interrupted the sentence “...participatory AI governance> through this, Ubuntuia can create...”.	Removed; the sentence now reads as intended.
§2.3	A stray hyphen and a doubled full stop interrupted two sentences: “... neutral technology. - it reflects...” and “...deepening existing inequalities..”.	Both corrected.
§3, §6, §8	Three section headings were followed by a lone hyphen on its own line.	Removed as drafting artifacts.
§8.4	The heading appeared as “7.4. Evaluation Timeline-”, carrying a section number inconsistent with its position and a trailing hyphen.	Renumbered to §8.4 and the trailing hyphen removed. The evaluation timeline is now presented as a table.
§7, Risk table	The header row began with an empty cell above the row numbers.	Labelled “#”.
Front matter	A “Contents” heading with no table of contents beneath it.	Removed as a drafting artifact.
Working group block	The first author’s surname was given as “Ngcob”, the fourth author’s name ran together as “Fadumoladan”, and the group administrator appeared as “Abdulkadir Fatima J”. All seven personal email addresses were included.	Names set as “Luvuyo Ngcobo” and “Fadumo Ladan Suleiman”, consistent with the surrounding text. The group mentor and group administrator are credited in the volume’s front matter rather than on the chapter opening, because several people who held those roles also served the volume as editors or mentors and are named once there. Emails withheld: publishing personal addresses in a permanent, publicly indexed DOI record is a privacy exposure; institutional affiliation is retained.

Location	As submitted	Editorial resolution
References	Eleven numbered entries, of which three were duplicates; three URLs were broken by stray spaces; several entries lacked locators.	Duplicates consolidated, URLs repaired and checked, DOIs added where they exist, and the list set alphabetically in APA 7th style. Documented in the editorial note above.

Education and Skills for the Future

National AI in Education Policy

A Comprehensive Policy Framework for Inclusive and Equitable AI-Enabled Education

CASE STUDY

The Republic of Ubutia

WORKING GROUP

Education and Skills for the Future Working Group

Abderrahim Chalfaouat · Hassan II University of Casablanca, Morocco

Abdiaziz Bashiir · DeepSom Institute of AI and Technology, Somaliland

Ahmed Esmat · The American University in Cairo, Egypt

Genet Abebe Gabre · Kotebe University of Education, Ethiopia

Hajer Ben Hadj Salem · High Institute of Human Sciences of Tunis, Tunisia

Nasiru Aliyu · Centre for Geodesy and Geodynamics, National Space Research and Development Agency, Toro, Nigeria

Naziru Mikail Abubakar · University of Westminster, London, United Kingdom

Nouruddeen Bashir Umar · Aliko Dangote University of Science and Technology, Wudil, Nigeria

Said Ali Said Yunus · State University of Zanzibar, Tanzania

Suggested citation: Chalfaouat, A., Bashiir, A., Esmat, A., Gabre, G. A., Ben Hadj Salem, H., Aliyu, N., Abubakar, N. M., Umar, N. B., & Yunus, S. A. S. (2026). *National AI in Education Policy*. In N. G. Abdulhamid (Ed.), *Eight Capstones, One Country* (pp. xx–xx). OpenSchool Initiative. <https://doi.org/10.5281/zenodo.22066741>

1. Executive Summary

The Republic of Ubuntia’s National AI in Education Policy (2026) marks a sovereign commitment to transforming our demographic energy — characterized by a median age of 20 and a 60% rural population — into a global competitive advantage. Rooted in the philosophy of Ubuntu, this framework transitions the national education system from a model of resource scarcity to one of AI-augmented abundance. Through leveraging adaptive learning to support overcrowded classrooms in ratios of 1:70 and deploying solar-powered, offline-capable AI hubs in borderlands, the policy ensures that high-quality, personalized education is no longer a privilege of the urban elite but a fundamental right for every learner from Lumora to the most remote agricultural districts.

Central to this transformation is the establishment of the National Educational Technology Accreditation Council (NETAC), which enforces strict standards for linguistic sovereignty and data protection. To prevent “epistemic colonialism,” all AI tools must operate natively in Ubuntu, English and French, while 30% of digital content must be locally co-created to reflect Ubuntian cultural values and oral traditions. This governance model mandates a “Human-in-the-Loop” architecture, ensuring that while AI automates administrative burdens for teachers, all high-stakes pedagogical and disciplinary decisions remain under human oversight to prevent algorithmic bias and protect student data sovereignty.

Implementation will follow a phased roadmap designed to catalyze a local innovation ecosystem and bridge the digital divide. Starting with teacher “AI Co-Pilot” certifications and the development of a national fibre backbone, the policy scales toward a complete curriculum pivot — moving from rote memorization to process-based assessment focused on critical thinking and AI literacy. Through strategic public-private partnerships and the incubation of local EdTech startups, Ubuntia aims to eliminate gender and rural disparities by 2030, positioning the nation as a continental leader in responsible, inclusive and ethical AI governance.

2. Context and Rationale

The Republic of Ubuntia is situated at a critical developmental inflection point, characterized by pronounced demographic dynamism and substantial institutional constraints. With a total population of 42 million and a median age of 20 years, the country exhibits a pronounced “demographic dividend” potential. However, this advantage is counterbalanced by an elevated youth unemployment rate exceeding 30%, with a disproportionate impact on women. Although gross enrollment rates at the primary level are high, severe attrition prior to completion of secondary education, combined with a highly selective and elite-oriented tertiary education system, risks excluding the majority of young people from the skills and qualifications required to participate effectively in an increasingly digital and globally integrated economy.

2.1 Structural Fragmentation

Ubuntia’s contemporary educational landscape is marked by pronounced structural fragmentation, most notably manifested in a substantial “digital divide” that separates technologically advanced urban innovation centres from rural districts, where more than 60% of the national population resides. The primary systemic challenges can be categorized into three interrelated domains: deficiencies in physical and digital infrastructure; disparities in classroom resources and learning environments; and a growing dependence on technology that is unevenly distributed and inequitably accessible across regions.

- National internet penetration stands at only 48%, with connectivity heavily concentrated in urban centres like Lumora and Kifaru City, leaving rural and borderland communities digitally isolated.
- Rural educators often manage overcrowded classrooms of over 70 pupils with nothing more than a chalkboard, while urban peers are already experimenting with AI-assisted tools and online platforms.
- Current AI pilots in agriculture, health and finance are largely donor-dependent and rely on imported systems that often fail to recognize local languages like Ubuntu or reflect national cultural heritage.

2.2 Technological Sovereignty as Policy Purpose

For Ubuntia, this policy is conceived not merely as a vehicle for technological adoption, but as an instrument for constructing a resilient and inclusive society through technological sovereignty. The integration of artificial intelligence into the education sector is anchored in the cultural value of Ubuntu (“humanity toward others”), thereby ensuring that technological advancement reinforces — rather than supplants — the indispensable roles of human mentorship and professional judgment that underpin high-quality education.

The policy is explicitly responsive to the urgency of the national context, prioritizing social cohesion, economic empowerment, cultural continuity, and ethical governance as core outcomes. It advances the deployment of solar-powered community hubs and offline-capable AI tools to provide personalized, adaptive instruction in rural classrooms, particularly in settings where individualized teaching by human educators is currently impracticable.

The framework stipulates that all educational technologies must be designed to reflect Ubuntia’s linguistic diversity and oral traditions, positioning AI as a means for the preservation and amplification of local cultural identities rather than as a conduit for “digital colonialism.”

In addition, the policy seeks to harness the demographic potential of youth for innovation by establishing localized AI skills-development pipelines, financed in part through diaspora remittances and integrated into emergent startup ecosystems in Kifaru City.

To ensure responsible implementation, the policy mandates the creation of national standards and regulatory oversight mechanisms aimed at identifying and mitigating algorithmic bias, protecting student data, and cultivating the delicate public trust necessary for a sustainable and equitable digital transformation.

Ubuntia cannot defer the initiation of its digital transformation until optimal connectivity conditions are achieved. In implementing a technology-agnostic, inclusion-oriented framework, AI is positioned as an enabling instrument for fostering social cohesion and promoting equitable socioeconomic development, thereby minimizing disparities in educational access for learners across both the remote mountainous interior and the densely populated coastal plains.

3. Problem Analysis

The Republic of Ubuntia faces educational stagnation amid rapid technological change, and the absence of a coordinated national strategy risks worsening inequality and undermining technological sovereignty.

3.1 The Rural–Urban Digital Divide

While cities like Lumora are starting to use digital platforms, about 60% of citizens live in rural agricultural areas with minimal infrastructure. As a result, national internet penetration is only 48%, leaving most rural and border communities without the connectivity needed for conventional AI tools. Rural schools also suffer from extreme overcrowding, with single teachers handling more than seventy pupils using only chalkboards, making personalized instruction unfeasible without automated, offline-capable systems.

3.2 Skills Gap and Educational Attrition

These structural challenges are worsened by a major skills gap and high educational attrition in a country where 60% of citizens are under 25. Despite this demographic potential, youth unemployment exceeds 30%, especially among women, because current curricula do not provide the future skills needed for a global AI-enabled economy. Although primary enrollment is high, attendance drops sharply before secondary completion, as the system offers few relevant or accessible higher-education pathways for non-elite students.

3.3 Donor Dependency and Digital Colonialism

Without a comprehensive national policy, Ubuntia has become a testing ground for donor-dependent, imported AI pilots that risk digital colonialism. Most tools are developed abroad and do not support the Ubuntu language or reflect local oral traditions and indigenous knowledge. These systems may reproduce historical inequities and biases that undermine Ubuntia’s constitutional commitment to equity. The lack of national interoperability standards also fragments systems and creates vendor lock-in, weakening state control over educational data and infrastructure.

3.4 The Trust Deficit

Finally, the deployment of large-scale technology in education faces a significant “trust deficit” fuelled by concerns over cybersecurity and professional autonomy. Without proper training, many teachers perceive AI as a threat to their pedagogical judgment or as a potential tool for surveillance. This environment of fragile public trust is further tested by the increasing reliance on digital platforms without robust data protection frameworks, leaving learners — particularly those in vulnerable borderland areas — exposed to exploitation and various cyber risks.

4. Policy Objectives

This policy aims to build a resilient, human-centred education system where AI promotes inclusion rather than inequality. By 2030, every school in Ubutia will have reliable connectivity, and every learner — regardless of gender, location, or ability — will be able to use AI to drive innovation and social mobility. To achieve this, the Republic of Ubutia commits to the following six strategic objectives:

1. To equip all learners, from primary to tertiary, and all educators with core digital skills, including critical thinking to detect algorithmic bias and a strong grasp of AI's ethical implications in a democratic society.
2. To bridge the rural–urban digital divide through innovative, technology-agnostic infrastructure — such as solar-powered community hubs and offline-capable AI tools — so that the 60% of the population in rural districts is not left behind.
3. To shift from consuming to creating technology by building national capacity to design and govern AI that reflects our cultural values and supports the Ubutia language.
4. To position educators as partners with responsibility by providing continuous professional development that transforms them into co-designers and ethical stewards of AI-enabled pedagogy, rather than passive recipients of automated tools.
5. To create a formal AI skills pipeline that reduces the 30% youth unemployment rate by preparing citizens for high-skilled roles in the global digital economy and supporting the growth of a local EdTech startup ecosystem.
6. To ensure public trust through the creation of a National AI Ethics Board and mandatory certification via the National Educational Technology Accreditation Council (NETAC), protecting learner data and preventing the deployment of biased or culturally misaligned systems.

5. Proposed Interventions

For the Republic of Ubuntuia to transition from vision to reality, it will need to implement a suite of targeted interventions designed to be mutually reinforcing. These actions ensure that technological progress remains balanced with social equity and the preservation of national identity.




Figure 2.1 The seven intervention pillars of the policy.

5.1 Infrastructure

With more than 60% of rural districts lacking stable electricity, the state will deploy decentralized solar-powered community learning hubs equipped with local servers. These modular units will host edge-computing AI tools that operate offline, delivering high-quality adaptive learning resources to borderland communities. In partnership with the Ministry of ICT and private telecommunications providers, the National Education Fiber Backbone (NEFB) will fast-track high-speed fibre connectivity to all secondary schools and vocational centres, targeting 80% national school connectivity by 2030 and prioritizing regions based on poverty indices, thereby establishing the digital floor.

5.2 Curriculum and Pedagogy

The Republic of Ubuntuia will shift from imported, culturally generic curricula to a framework rooted in local linguistics and social realities. All AI-enabled educational software must support the Ubuntu language, while a national project digitizes oral storytelling and indigenous knowledge for integration into AI training datasets, safeguarding local wisdom from global algorithms. To accommodate areas without screens, the primary curriculum will incorporate unplugged AI modules featuring storytelling, logic games and analog problem-solving to build foundational algorithmic thinking and data ethics skills.

5.3 Teacher Empowerment

Teachers serve as the primary practitioners in the education system; accordingly, interventions emphasize support over replacement. The National Teacher Digital Residency (NTDR) mandates certification in AI-assisted pedagogy for every educator, enabling AI to handle administrative tasks like grading and lesson planning so that teachers can focus on mentoring and emotional guidance. Digital and physical Professional Learning Communities (PLCs) will connect teachers to share “Made-in-Ubuntia” best practices, promoting collaborative innovation instead of top-down mandates.

5.4 Economic Pipeline

To leverage the demographic dividend, education will connect directly to global and local digital economies. The government will offer seed funding and tax incentives for the Kifaru City EdTech Incubator, empowering local startups to create AI solutions tailored to Ubuntian contexts and positioning the nation as an exporter of African innovation. Vocational “fast-tracks” will provide six-month certifications in data labelling, AI maintenance and digital entrepreneurship at centres targeting non-elite youth, especially young women and rural migrants, for immediate access to high-skilled jobs to address the 30% youth unemployment rate.

5.5 Governance

Public trust requires rigorous oversight to remain intact. The National Educational Technology Accreditation Council (NETAC) will certify all AI tools for deployment in schools, evaluating pedagogical effectiveness, interoperability to avoid vendor lock-in, and cultural sensitivity. The National AI Ethics Board (NAIB) will audit systems for bias, while a Data Sovereignty Protocol (DSP) will keep all student data within national borders under robust encryption, restricted solely to educational improvements and barring commercial surveillance.

5.6 Inclusive Special Education and Accessibility

Assistive AI deployment will prioritize technologies like real-time speech-to-text and text-to-sign language for learners with disabilities. These tools will be distributed through a national equity grant, ensuring children with special needs lead the benefits of the AI revolution.

5.7 Public–Private–Community Partnerships

The “Ubuntu” Innovation Fund will pool resources from government, diaspora and international partners to issue micro-grants for community-led digital projects. This approach transforms the policy from a mere government document into a people-owned national movement.

These seven pillars form a holistic strategy for national development, emphasizing rural access, linguistic sovereignty and teacher empowerment to position AI as a great equalizer. In turn, Ubuntia’s youthful population will emerge as a global powerhouse of responsible innovation.

6. Governance and Implementation Model

The Republic of Ubuntuia’s shift to an AI-enabled education system requires a multi-sectoral governance framework that balances innovation and cultural integrity, ensuring technology supports social equity and national sovereignty. This dynamic, decentralized model replaces centralized power by empowering independent statutory bodies to manage technological risks while enhancing local innovation under human-centred oversight.

6.1 Institutional Oversight

The governance model is structured around two statutory bodies established to mediate between central ministries and resource-constrained local governments.

Table 2.1 Statutory oversight bodies.

Body	Mandate
National Educational Technology Accreditation Council (NETAC)	The technical gatekeeper. NETAC is responsible for certifying all AI-enabled tools used in schools, evaluating software for pedagogical effectiveness, technical interoperability (to prevent vendor lock-in), and linguistic compatibility with the Ubuntu language.
National AI Ethics Board (NAIB)	An independent body that systematically audits algorithms to identify and address socio-cultural biases, ensures compliance with constitutionally protected rights and the Data Protection Act of 2021, and supervises normative, legal and ethical standards for AI deployment.

6.2 Data Sovereignty and Ethics

To mitigate the risks associated with digital colonialism, the state will implement a National Data Sovereignty Protocol.

- **Local data hosting.** All learner-related data shall be stored exclusively on domestically located servers in order to maintain national and institutional control over critical educational digital infrastructure.
- **Data protection and privacy safeguards.** Building on the 2021 Data Protection Act, strict standards for data encryption, anonymization, and secure access management will be mandated to reduce cyber exploitation, with particular emphasis on protecting vulnerable populations in rural and border regions.

6.3 Sustainable Financing Strategy

Implementation will be sustained through diversified financing that leverages Ubuntuia’s macroeconomic and sectoral strengths. The Ubuntu Innovation Fund will serve as a central mechanism, pooling government budget allocations, international development grants, and “Diaspora Digital Bonds” backed by annual remittances. Public–Private–Community Partnerships (PPCP) will be institutionalized via fiscal incentives for the Kifaru City technology hub to develop “made-in-Ubuntia” digital solutions, and through zero-rated data access agreements with telecommunications operators for all accredited educational platforms.

6.4 Strategic Implementation Roadmap

The rollout follows a phased approach to manage institutional capacity and build public trust.




Figure 2.2 Three-phase implementation roadmap from Year 1 to Year 10.

Table 2.2 Strategic implementation roadmap.

Phase	Period	Focus	Core actions
Phase I	Years 1–2	Foundation	Formalize NETAC and NAIB; launch “Solar Hub” pilots in rural districts; initiate the National Teacher Digital Residency.
Phase II	Years 3–5	Scaling	Expand AI-integrated curricula to 50% of secondary schools; mandate Ubuntu-language support for all licensed software.
Phase III	Years 6–10	Maturity	Achieve universal school connectivity; evaluate the impact on reducing the 30% youth unemployment rate through AI vocational pathways.

7. Risk and Mitigation Plan

The transition toward AI-enabled education in the Republic of Ubuntu entails substantial systemic risks. The table below summarises the risks identified and the proposed mitigation strategies to safeguard Ubuntu's 42 million inhabitants and to ensure the long-term viability of the educational reform agenda.

Table 2.3 Risk mitigation plan.

Risk Type	Identified Risk	Mitigation Strategy
Digital Colonialism	Over-reliance on donor-funded, imported AI systems that lack cultural fit and erode local identity.	Mandate NETAC certification for all tools to ensure linguistic support for Ubuntu and local cultural contextualization.
Widening Inequality	AI benefits being restricted to urban elites in Lumora, deepening the divide for the 60% rural population.	Deploy solar-powered hubs and offline-capable AI modules designed for low-connectivity, overcrowded rural classrooms.
Data Exploitation	Cyber fraud and unauthorized commercial use of sensitive student data due to weak enforcement capacity.	Implement the National Data Sovereignty Protocol, requiring domestic data hosting and strict anonymization under the 2021 Data Protection Act.
Algorithmic Bias	AI systems replicating historical inequities or discriminatory patterns against marginalized groups and women.	Establish the National AI Ethics Board (NAIB) to perform independent, systematic audits of all educational algorithms before deployment.
Teacher Resistance	Fear of job displacement or increased administrative burden leading to low adoption of digital tools.	Position teachers as “custodians of trust” through continuous professional development that focuses on reducing workload while preserving pedagogical autonomy.

8. Monitoring and Evaluation Framework

To ensure the successful delivery of the AI in Education strategy, the Republic of Ubutia will implement a multi-tiered monitoring and evaluation system that tracks progress from rural agricultural districts to the technology hubs of Kifaru City.

8.1 Key Performance Indicators for 2030

Success will be measured against five strategic benchmarks. Each benchmark derives from one of the intervention pillars set out in §5, as shown in the final column.

Table 2.4 Key performance indicators for 2030, mapped to their source intervention pillars.

Benchmark	Target for 2030	Source pillar
Infrastructure Equity	Achieve 80% high-speed connectivity in rural schools and borderland frontier districts, reducing the current urban–rural digital divide by half.	Pillar 1 Infrastructure
Teacher Empowerment	100% of the teaching workforce certified under the National Teacher Digital Residency programme, with a 30% reported reduction in administrative burden via AI automation.	Pillar 3 Teacher Empowerment
Economic Impact	A measurable 15% decrease in the 30% youth unemployment rate through the successful placement of graduates in AI-vocational and digital service sectors.	Pillar 4 Economic Pipeline
Linguistic Integrity	100% of nationally accredited AI educational software to feature full instructional and interface support in the Ubutia language.	Pillar 2 Curriculum and Pedagogy
Data Sovereignty	Zero instances of unauthorized cross-border student data transfer, verified by annual independent audits of the National AI Ethics Board (NAIB).	Pillar 5 Governance

Pillars 6 (Inclusive Special Education and Accessibility) and 7 (Public–Private–Community Partnerships) are delivered through the evaluation mechanisms in §8.2 rather than through a headline 2030 benchmark: accessibility provision is tracked by the National AI Education Dashboard, and community partnership outcomes are captured through the district-level Professional Learning Communities.

8.2 Evaluation Mechanisms and Oversight

To maintain transparency and public trust, the following mechanisms will be institutionalized:

- **The National AI Education Dashboard.** A real-time, public-facing platform managed by NETAC to track device distribution, solar-hub functionality, and learner progress across all 42 million citizens.
- **Bi-annual ethical audits.** The NAIB will conduct mandatory reviews of algorithmic performance to identify and rectify socio-cultural biases, ensuring AI tools do not disadvantage marginalized ethnic groups or women.
- **Community feedback loops.** Professional Learning Communities at district level, where rural teachers and parents can provide qualitative feedback on tool usability and cultural fit.

8.3 Continuous Improvement Cycle

This framework is a “living” system. Every three years, a strategic policy review will be conducted to assess the impact of AI on Ubuntu’s social mobility. Results will be used to recalibrate Ubuntu Innovation Fund allocations, ensuring that financial resources are directed where the last mile of implementation is most challenged.

9. Conclusion and Call to Action

The Republic of Ubutia stands at a historic turning point. Integrating artificial intelligence into education is not just a technological upgrade but an act of national self-determination. Rooted in the philosophy of Ubuntu, Ubutia rejects the passive adoption of foreign tools and instead chooses technology that amplifies our humanity. The policy intends to move beyond 1:70 classroom ratios and 30% youth unemployment toward AI-augmented abundance, where a child in the most remote border district has the same innovative potential as one in the heart of Lumora.

To transform this policy into a lived reality for our 42 million citizens, the following stakeholders are called to mobilize with immediate urgency:

- **Ubutia’s educators.** You are “custodians of trust.” The state will continually empower you through the National Teacher Digital Residency, ensuring AI serves you — not the other way around.
- **Private sector and diaspora.** Invest in the Ubuntu Innovation Fund to build a homegrown EdTech ecosystem in Kifaru City and export Ubutian ingenuity across the continent.
- **International partners.** Join us in peer-to-peer partnerships that respect our data sovereignty, our linguistic integrity, and our right to local co-creation.
- **Ubutia’s youth.** Prepare to lead. The digital floor has been set; the ceiling of your innovation is now limited only by your imagination.

The time for pilot projects has passed; the era of systemic transformation has begun. Let us act with the collective courage that our “Unity in Diversity” demands. By 2030, the Republic of Ubutia will not just be a participant in the digital age, but a sovereign leader of the AI revolution, ensuring that in the march of progress, no learner is left behind.

References

The submitted chapter did not include a reference list. The volume editors have identified authoritative sources for the external frameworks and scholarship the chapter engages. Every URL and DOI below was checked and resolved successfully at the time of publication.

- African Union. (2024). *Continental artificial intelligence strategy: Harnessing AI for Africa's development and prosperity*. African Union.
<https://au.int/en/documents/20240809/continental-artificial-intelligence-strategy>
- Birhane, A. (2020). Algorithmic colonization of Africa. *SCRIPTed*, 17(2), 389–409.
<https://doi.org/10.2966/scrip.170220.389>
- International Telecommunication Union. (2024). *Measuring digital development: Facts and figures 2024*. ITU. <https://www.itu.int/itu-d/reports/statistics/facts-figures-2024/>
- Miao, F., & Holmes, W. (2023). *Guidance for generative AI in education and research*. UNESCO.
<https://doi.org/10.54675/EWZM9535>
- UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. UNESCO.
<https://unesdoc.unesco.org/ark:/48223/pf0000381137>
- UNESCO. (2023). *Global education monitoring report 2023: Technology in education — A tool on whose terms?* UNESCO. <https://doi.org/10.54676/UZQV8501>
- United Nations. (2015). *Transforming our world: The 2030 Agenda for Sustainable Development* (A/RES/70/1). United Nations. <https://sdgs.un.org/2030agenda>

Editorial note on sources

No references were submitted with this chapter. The sources above were located by the volume editors to support the frameworks the chapter invokes: the UNESCO Recommendation on AI ethics and the UNESCO guidance on generative AI in education, which underpin the NETAC and NAIB certification model; the UNESCO Global Education Monitoring Report on technology in education, the standard reference for the equity argument the chapter advances; the African Union Continental AI Strategy, the regional instrument most directly relevant to the sovereignty argument; the ITU digital development statistics, the standard source for connectivity baselines; Birhane (2020) on algorithmic colonization, which corresponds closely to the chapter's "epistemic colonialism" and "digital colonialism" framing; and the 2030 Agenda. They are offered as supporting apparatus; the authors' analysis and wording are unchanged.

Verification of claims

Table 2.5 Editorial verification of claims.

Claim as written	Section	Editorial verification
“A total population of 42 million and a median age of 20 years”	§2	Internal to the case study, and consistent with the figures used across other chapters in this volume.
“National internet penetration stands at only 48%”	§2.1, §3.1	Internal to the case study, and consistent with Thematic Areas 4, 7 and 8.
“More than 60% of the national population resides in rural districts”	§2.1	Internal to the case study, and consistent with Thematic Areas 6, 7 and 8.
“Youth unemployment exceeding 30%”	§2, §3.2	Internal to the case study, and consistent with Thematic Areas 6 and 8.
“Classrooms of over 70 pupils” and “1:70 classroom ratios”	§2.1, §3.1	Internal to the case study. No external source is cited.
“The Data Protection Act of 2021”	§6.1, §6.2, §7	Internal to the case study. The same instrument is referenced in Thematic Area 7, indicating a shared scenario parameter rather than a real statute.
All 2030 targets (80% connectivity, 100% teacher certification, 15% unemployment reduction, 30% locally co-created content)	§1, §5.1, §8.1	Authors’ own policy targets for the case study, not empirical findings.
“Epistemic colonialism” and “digital colonialism”	§1, §2.2, §3.3	Externally supported as a concept. Birhane (2020) sets out the algorithmic colonization argument the chapter advances. Citation supplied by the editors.

Editorial judgements

No author review cycle was available before publication. Where the submitted text contained defects whose intended meaning could be established with confidence from the surrounding argument, the volume editors have corrected them rather than leaving them in the published record. Each such correction is listed below. No substantive claim, target, or policy position has been altered.

Table 2.6 Editorial judgements applied to this chapter.

Location	As submitted	Editorial resolution
§9, opening paragraph	“Rooted in the philosophy of Ubuntu, adopting foreign tools and instead choose technology that amplifies our humanity is rejected passively.”	The sentence had lost its subject and main verb in drafting, though all its elements were present. Restored to: “Rooted in the philosophy of Ubuntu, Ubuntu rejects the passive adoption of foreign tools and instead chooses technology that amplifies our humanity.” The meaning is the authors’ own and is stated plainly elsewhere in §2.2 and §3.3.
§9, opening paragraph	“The policy intends to actualize moving beyond 1:70 classroom ratios...”	Corrected to “The policy intends to move beyond 1:70 classroom ratios...”.
§5 and §8.1	The chapter refers to “these seven pillars” but then measures success against “five strategic benchmarks derived from the country’s core policy pillars,” without stating the relationship.	Each of the five benchmarks maps unambiguously to one intervention pillar, so the mapping has been made explicit in Table 2.5. Pillars 6 and 7 carry no headline 2030 benchmark; their delivery is tracked through the evaluation mechanisms the authors set out in §8.2, and this is now stated.
Title block	“The Republic of UbuntuNational AI in Education Policy”	Missing separator restored.
§3.1	“more than seventy (70 pupils” — unclosed parenthesis	Redundant parenthetical removed; reads “more than seventy pupils”.
§4, objective 4	“To position educators as with responsibility”	Restored to “To position educators as partners with responsibility”.
§5.3	“Teachers serve as the primary practitioner in the education system”	Number agreement corrected to “practitioners”.
§8.1	“derived from the countries core policy pillars”	Possessive corrected to “the country’s core policy pillars”.
Front matter	A “Contents” heading with no table of contents beneath it.	Removed as a drafting artifact. Navigation is provided by the volume’s own table of contents.

Labour, Employment and Future of Work

AI and the Future of Work Policy Blueprint

Protect Today's Workers While Preparing Ubuntu's Youth for Tomorrow's Digital Economy

CASE STUDY

The Republic of Ubuntu

WORKING GROUP

Labour, Employment and the Future of Work Working Group

Shuaibu Sambo Othman · Voice of Nigeria, Nigeria

Hadiza N. Tijjani · Nigerian Television Authority, Nigeria

Lorna Okatch · Government of Kenya, Kenya

Sarah Ibrahim · Cairo University, Egypt

Aisha Usman Muhammad · Gombe State University, Nigeria

Suggested citation: Othman, S. S., Tijjani, H. N., Okatch, L., Ibrahim, S., & Muhammad, A. U. (2026). AI and the Future of Work Policy Blueprint. In N. G. Abdulhamid (Ed.), *Eight Capstones, One Country* (pp. xx–xx). OpenSchool Initiative.
<https://doi.org/10.5281/zenodo.22066741>

1. Executive Summary

This Artificial Intelligence Policy Blueprint for Labour, Employment, and the Future of Work is tailored to the African context, with specific application to the Ubuntia Nation. It responds to the accelerating adoption of AI-driven technologies across public and private sectors and their profound implications for employment structures, workforce skills, productivity, and social protection systems.

For Ubuntia, AI represents both an opportunity to leapfrog structural development constraints and a risk of deepening unemployment, informality, and inequality if left unguided. This blueprint provides a coherent national framework to ensure AI adoption supports decent work, inclusive growth, and socio-economic resilience. It prioritises human-centred AI, protection of workers' rights, gender and youth inclusion, and alignment with African Union digital transformation agendas.

The policy outlines strategic interventions across skills development, labour regulation, innovation, governance, and monitoring. Implemented effectively, it will position Ubuntia as a continental leader in responsible AI governance for the future of work.

1.1 Policy Focus

- National reskilling initiatives
- AI adoption in public service
- Worker protections in the gig economy
- Incentives for inclusive job creation

2. Context and Rationale

Africa has the world's youngest population and a rapidly expanding labour force. In Ubutia, over 60% of workers are engaged in informal employment, with high levels of youth and graduate unemployment. At the same time, AI technologies are increasingly deployed in finance, agriculture, public administration, recruitment, and platform-based services.

The convergence of demographic pressure, digitalisation, and global competition makes AI governance in labour markets a strategic necessity. Without policy direction, AI risks exacerbating job displacement, digital exclusion, and labour precarity. With the right framework, however, AI can improve productivity, formalise informal work, enhance service delivery, and connect Ubutian workers to regional and global value chains.

This blueprint aligns with the African Union Digital Transformation Strategy (2020–2030), Agenda 2063, and international labour standards, while responding to Ubutia's national development priorities.

3. Problem Analysis

“The Republic of Ubuntuia is facing major development challenges at a time when digital technologies and artificial intelligence are spreading quickly across the world” (Country Profile: The Republic of Ubuntuia, n.d.).

3.1 Youth Unemployment Meets Technological Disruption

Ubuntuia faces a stark disconnect between rapid AI innovation and its labour market’s capacity for equitable transition, risking deepened inequalities for vulnerable groups.

Youth unemployment exceeds 30%, with young women hit hardest in informal sectors vulnerable to AI displacement. High primary enrollment belies early dropouts and limited tertiary access, leaving youth without essential digital skills. Small-scale farmers grapple with climate variability and AI tool inaccessibility, while fintech workers endure cyber fraud amid weak protections, as seen in Lumora’s near-collapse.

3.2 The Digital Divide and Algorithmic Bias

AI readiness varies sharply by geography, gender, and income: urban workers with connectivity seize opportunities, but rural, female, and low-income groups lag due to poor infrastructure and devices. Imported AI often lacks local languages and datasets, yielding biased hiring, credit scoring, and services that entrench marginalization.

3.3 Weak Oversight in a Rapidly Digitalizing Economy

Regulations trail AI’s labour impacts. The Data Protection Act sets data rights, but enforcement falters from corruption, under-resourced courts (especially outside the capital), and lax oversight of algorithms, platform work, and data practices — enabling unchecked deployments that erode trust and burden workers.

3.4 Digital Sovereignty and Inclusive Development

Without intervention, how can Ubuntuia curb AI from worsening gender inequalities, sidelining women’s roles, and letting foreign technology reshape labour markets against inclusive, decent Ubuntuian work? The imperative: claim digital sovereignty and bolster institutions for AI-driven equitable growth.

3.5 Human-Centred AI Across Key Sectors

Rather than replacing labour, AI in Ubuntuia must augment human work:

- **Agriculture.** AI should support farmers through weather alerts, pest warnings, and price information via low-tech channels, not displace farm labour.
- **Health.** AI should assist triage, records, and disease monitoring to reduce workload, not replace scarce health workers.
- **Education.** AI should function as a teaching assistant, not a substitute for teachers, particularly in overcrowded rural classrooms.

3.6 AI-Specific Risks

- Algorithmic bias in recruitment and performance management
- Worker surveillance and data privacy violations

- Job polarisation between high-skill and low-skill workers

Table 3.1 Vulnerability of employment sectors in Ubuntia.

Sector	AI Adoption Level	Employment Risk	Policy Priority
Agriculture	Medium	Medium	High
Public Service	Medium	Low	Medium
Financial Services	High	High	High
Informal Trade	Low	Medium	High
Routine Clerical Jobs	High	High	High
Technical Skills Roles	Medium	High	Medium

4. Policy Objectives

Table 3.2 The five policy objectives.

Objective	Commitment
Developing Future-Ready Digital Skills for Citizens	Equip citizens with digital literacy, technical training, and lifelong learning opportunities integrated into school curricula and community learning centres.
Promoting Equitable AI Adoption in Public Services	Ensure inclusive, ethical AI deployment in healthcare, education, agriculture, and governance, with safeguards for vulnerable groups and persons with disabilities.
Protecting Gig Workers through Legal and Social Frameworks	Recognize gig workers legally, guarantee fair remuneration, provide social security, and establish dispute resolution mechanisms.
Stimulating Inclusive Job Creation in Rural and Urban Areas	Support digital entrepreneurship, innovation hubs, SME funding, and remote work opportunities to balance rural and urban growth.
Strengthening Digital Infrastructure for Innovation	Invest in broadband, cybersecurity, secure data systems, and public–private partnerships to build a resilient digital ecosystem.

5. Proposed Interventions

Six interventions translate the objectives into deliverable programmes. Each names an implementing body and measurable indicators.

5.1 Community AI Work-and-Learn Hubs (CAWL)

Table 3.3 Community AI Work-and-Learn Hubs (CAWL).

Element	Detail
Description	Establish hubs offering micro-credentials in AI-assisted work, integrated into school curricula and community learning centres. Apprenticeships connect trainees to schools, clinics, cooperatives, and startups.
Implementing Body	Ministry of Education and TVET; Ministry of Labour; local governments; women's associations
Indicators	• 30,000 trainees • ≥60% women participation • ≥50% job placement • Integration into national curricula

5.2 Ubuntu-First Agri-Advisory AI

Table 3.4 Ubuntu-First Agri-Advisory AI.

Element	Detail
Description	Provide weather, pest, and price information via USSD, SMS and voice in local languages. Designed for low-bandwidth environments and inclusive of women farmers.
Implementing Body	Ministry of Agriculture; farmer cooperatives; telecommunications operators
Indicators	• ≥1 million farmers reached • ≥40% women farmers • ≥10% reduction in post-harvest loss

5.3 Health Workflow Augmentation

Table 3.5 Health Workflow Augmentation.

Element	Detail
Description	AI tools for triage, records, and disease monitoring to support — not replace — frontline staff. Accessibility guidelines ensure inclusion of persons with disabilities.
Implementing Body	Ministry of Health; district health offices; universities
Indicators	• ≥25% reduction in clinic wait times • ≥90% record completeness • ≥70% of staff report workload relief

5.4 Fair Work Code for Digital Platforms

Table 3.6 Fair Work Code for Digital Platforms.

Element	Detail
Description	A governance package ensuring algorithmic transparency, fair pay, grievance redress, and portable benefits. Legal recognition of gig workers and linkage to social security systems.
Implementing Body	Ministry of ICT and Innovation; Ministry of Labour; Data Protection Authority
Indicators	• ≥80% of platforms certified • ≥20% increase in dispute resolution • ≥50% of workers enrolled in portable benefits • Legal frameworks enacted

5.5 Inclusive Job Creation Incentives

Table 3.7 Inclusive Job Creation Incentives.

Element	Detail
Description	Provide targeted SME funding, youth startup support, and remote work opportunities. Incentives for rural innovation hubs to curb urban migration.
Implementing Body	Ministry of Labour; Ministry of ICT; local governments; diaspora investors
Indicators	• ≥500 SMEs funded • ≥200 youth startups supported • ≥30% increase in rural employment opportunities

5.6 Ubuntu AI Commons and Digital Infrastructure

Table 3.8 Ubuntu AI Commons and Digital Infrastructure.

Element	Detail
Description	Build sovereign datasets and language resources; update procurement rules to favour local solutions. Invest in broadband expansion, cybersecurity, and public–private partnerships.
Implementing Body	National Statistics Office; Ministry of ICT; telecommunications operators; universities
Indicators	• ≥50 datasets released • ≥3 Ubuntu language models • ≥30 public projects using commons assets • ≥20% increase in broadband coverage

6. Governance and Implementation Model

Effective implementation of the AI and the Future of Work Policy Blueprint in Ubuntuia requires a coordinated, transparent, and multi-level governance framework that aligns national strategy with local realities, safeguards public trust, and ensures that AI adoption remains human-centred and inclusive.

6.1 Institutional Governance Framework

The Ministry of Labour, Skills and Social Transformation (MLSST) shall serve as the lead coordinating authority for the blueprint. It will provide overall policy oversight, ensure alignment with labour standards, workforce transition strategies, and social protection reforms, and coordinate inter-ministerial implementation.

To support strategic direction and accountability, a National AI and Future of Work Council (NAFWC) shall be established as a high-level, multi-stakeholder body.




Figure 3.1 Governance framework for AI and the future of work, from strategic oversight to local implementation.

Table 3.9 National AI and Future of Work Council: composition and roles.

Stakeholder Group	Key Role
MLSST (Chair)	Policy leadership and labour oversight
Digital Economy and Innovation	AI strategy and digital infrastructure
Education and Human Capital	Curriculum reform and skills development
Agriculture and Rural Development	AI for smallholder productivity
Health	AI-assisted service delivery
Finance and National Planning	Funding and fiscal alignment
Data Protection Authority	Ethical AI and data governance
Local Governments	Community-level implementation
Employers' Associations	Private sector adoption
Trade Unions and Worker Groups	Worker representation
Universities and Research Institutions	Evidence and innovation
Civil Society, Youth and Women's Groups	Inclusion and accountability
Technology and Platform Firms (advisory)	Technical expertise

Core functions of the Council

- Approve national priorities for AI deployment in labour-intensive sectors
- Ensure alignment between AI innovation, labour protection, and skills policy
- Oversee ethical compliance and human-centred AI principles
- Facilitate social dialogue and conflict resolution
- Endorse annual implementation and evaluation reports

6.2 Sectoral and Local Implementation Architecture

**Figure 3.2** Digital access and decentralised AI implementation zones in Ubuntia.

The map illustrates disparities in digital access between urban innovation hubs and rural districts, underscoring the need for a decentralised, local-government-led approach to implementing AI tools that support workers across sectors.

Implementation shall follow a sector-specific but coordinated approach, recognizing that AI affects agriculture, health, education, and public services differently. Each priority sector shall

establish a Sector AI Focal Unit within the relevant ministry to translate national policy into sector-appropriate, job-supporting AI solutions.

Table 3.10 Sector AI Focal Units and responsibilities.

Sector	Lead Ministry	Core Responsibilities
Agriculture	Agriculture and Rural Development	Farmer-support AI tools, extension systems
Health	Ministry of Health	AI-assisted triage, surveillance, records
Education	Education and Human Capital	Teacher-support tools, learning platforms
Public Services and Youth Employment	Digital Economy and Innovation	Skills platforms, local service planning

These units will identify AI use cases that support rather than replace workers, coordinate pilot projects and scaling, ensure tools are accessible, low-bandwidth, and locally relevant, and work closely with unions, professional bodies, and frontline workers.

Local governments will play a critical role in implementation by hosting community digital and AI support hubs, supporting farmer cooperatives, schools, and clinics, facilitating training for extension officers, teachers, and health workers, and providing feedback on labour impacts and inclusion outcomes.

6.3 Regulatory, Oversight, and Integrity Mechanisms

The National Data Protection Authority, in collaboration with MLSST, shall oversee compliance with the Data Protection Act (2021), transparency and accountability of AI systems used in employment, and protection against algorithmic discrimination and intrusive worker surveillance. High-risk AI systems used in recruitment, performance management, credit scoring, or platform work shall be subject to registration and disclosure requirements, human-in-the-loop safeguards, and periodic audits for bias and labour impact.

To address historical challenges of corruption and public mistrust, AI procurement shall follow open, competitive, and traceable processes, supported by independent audits of publicly funded AI projects. Civil society and parliamentary committees shall be empowered to review outcomes and ensure accountability.

6.4 Financing, Capacity Building, and Accountability

Implementation will be supported through a blended financing model, including national budget allocations, employer co-financing for reskilling, innovation and digital transformation funds, diaspora investment channels, and development partner grants and technical assistance. Public-private partnerships will be encouraged, provided they align with human-centred AI principles, local capacity building, and job protection and skills transfer commitments.

A National Observatory on AI, Labour, and Skills shall be established to track labour market impacts, support skills forecasting and curriculum reform, generate evidence for policy adjustment, and document best practices. Universities, technical institutes, and research centres will serve as key partners.

Accountability will be ensured through annual AI and Future of Work implementation reports submitted to Cabinet and Parliament, public dashboards on employment, skills, and inclusion indicators, and annual tripartite consultations to review progress and adjust strategy.




Figure 3.3 AI policy implementation pathway, from policy to worker-level outcomes.

The governance and implementation model ensures that AI adoption in Ubuntu is coordinated, transparent, inclusive, and accountable, enabling productivity and service delivery gains while protecting jobs, supporting workers, and reinforcing public trust.

7. Risk and Mitigation Plan

Introducing AI into Ubuntuia's labour market is not just about efficiency; it brings specific dangers that could widen existing gaps if we are not careful. Based on our vulnerability analysis, we anticipate three major risk areas and have designed specific guardrails for each.

Table 3.11 Risks and mitigation measures.

Risk	Description	Mitigation
Risk 1 The “black box” trap in hiring and management	There is a high risk that automated recruitment tools, often imported from abroad, will filter out capable Ubuntuian youth because their CVs do not look like those in Europe or the United States, or because they lack formal degrees despite having practical skills. Furthermore, algorithmic management in the gig economy could penalize workers unfairly without human recourse.	Enforce a human-in-the-loop mandate for all high-stakes employment decisions — hiring, firing, and disciplinary actions. Companies deploying these tools must pass a local bias audit, proving their algorithms can interpret Ubuntuian educational and work contexts correctly.
Risk 2 The digital gender divide deepening	Women in Ubuntuia already face steeper barriers in the workforce, often relying on informal micro-enterprises. If AI tools are designed solely for smartphone users in Kifaru City with high-speed 4G, they will exclude the female farmer in the Central Highlands or the trader using USSD.	The “Ubuntu-First” standard. Publicly funded AI labour tools must work on low-bandwidth devices and support voice interaction in Ubuntu. Gender impact assessments will be required before rolling out any government automated service.
Risk 3 Data extraction without value return	There is a danger that Ubuntuia becomes a “data mine” where the workforce generates data for foreign platforms, but the value — and the intellectual property — leaves the country.	Implement data sovereignty clauses in public procurement. If a foreign firm uses Ubuntuian workforce data to train its models, it must share the resulting insights or derived models with the National AI Council to benefit the local ecosystem.

8. Monitoring and Evaluation Framework

To ensure this policy actually leads to decent work and not just “digital buzz,” we are moving beyond vanity metrics — such as the number of apps downloaded — to measure actual labour outcomes. This framework relies on a partnership between the Ministry of Labour, trade unions, and civil society.

8.1 The “Living Lab” Approach

We cannot wait two years to see if a policy is working. We will treat Kifaru City’s technology sector and the northern rural districts as two distinct “living labs”.

Quarterly check-ins. The National AI and Future of Work Council will hold quarterly reviews with gig worker representatives and employer federations to catch unintended consequences early, such as sudden drops in gig wages or spikes in algorithmic account suspensions.

8.2 Key Performance Indicators

Success will be tracked across four dimensions.

Table 3.12 Key performance indicators by dimension.

Dimension	Indicator	Target (Year 3)	Source of Data
Worker Protection	Resolution time: average time to resolve a dispute between a gig worker and an algorithmic platform	Under 48 hours	Ministry of Labour Digital Desk
Skills and Mobility	Placement rate: percentage of graduates from Community AI Hubs who secure a job or increase freelance income within 6 months	Over 50%	TVET tracking system
Inclusion	Gender ratio: percentage of women accessing AI-assisted agricultural or business tools	Over 40%	Cooperative audits
Fairness	Algorithm certification: percentage of major digital labour platforms that have passed the national Fair Work audit	Over 80%	Data Protection Authority

8.3 Qualitative Feedback: “Ground Truth” Stories

Numbers hide nuances. To capture the reality of the rural teacher managing 70 pupils or the entrepreneur facing cyber fraud, we will conduct:

- **Annual worker voice surveys.** Anonymous feedback channels via SMS and USSD for workers to report how AI tools are impacting their stress levels and autonomy.
- **Shadow audits.** “Mystery shopper” exercises where ministry officials apply for jobs or work on platforms to test whether the anti-bias and transparency regulations are actually being followed in practice.

8.4 Adaptive Policy Review

This blueprint is not written in stone. An annual “State of Future Work” report will be presented to Parliament. If specific interventions — for example the SME grants — fail to produce jobs for two consecutive cycles, the funding will be automatically triggered for review

and reallocation to more effective mechanisms. This ensures we remain agile enough to keep up with the technology while keeping our promise to the workers of Ubuntia.

9. Conclusion and Call to Action

Artificial intelligence is reshaping the foundations of labour, employment, and the future of work. For Ubuntuia, the challenge is not whether AI will transform the economy, but how it will be harnessed to serve the people equitably. This blueprint has outlined the risks of unchecked adoption — algorithmic bias, job displacement, and digital exclusion — while also highlighting the immense opportunities for productivity, inclusion, and innovation.

The path forward requires decisive leadership, inclusive governance, and sustained investment in human capital. AI must be deployed as a tool to empower workers, not replace them; to expand opportunities, not deepen inequalities. By embedding fairness, transparency, and accountability into AI systems, Ubuntuia can set a continental precedent for responsible innovation.

9.1 We Call Upon

- **Government.** To adopt and enforce this framework, ensuring labour laws, social protections, and digital infrastructure evolve in step with AI.
- **The private sector.** To commit to ethical AI practices, invest in reskilling, and prioritize job-creating innovations.
- **Academia and civil society.** To champion research, advocacy, and worker voice in shaping AI's role in employment.
- **Workers and unions.** To engage actively in dialogue, ensuring their rights and perspectives remain central to AI governance.
- **Regional and international partners.** To support Ubuntuia's leadership in building an inclusive AI-driven labour market aligned with African Union and global standards.

The future of work in Ubuntuia will not be written by algorithms alone; it will be determined by the collective choices we make today. This is a call to action: to embrace AI with courage, to govern it with wisdom, and to ensure that every worker, regardless of gender, age, or background, shares in the prosperity of the digital era.

References

This chapter was submitted with a reference list. The volume editors have normalised it to APA 7th style, restored one entry that had been split across two lines, and verified every locator. Where the authors supplied no URL or DOI, one has been added.

- African Union. (2020). *The digital transformation strategy for Africa (2020–2030)*. African Union Commission. <https://au.int/en/documents/20200518/digital-transformation-strategy-africa-2020-2030>
- African Union Commission. (2024). *Continental artificial intelligence strategy*. Adopted at the 45th Ordinary Session of the Executive Council, Accra, Ghana. African Union. <https://au.int/en/documents/20240809/continental-artificial-intelligence-strategy>
- Country profile: The Republic of Ubuntu*. (n.d.). Fellowship case document. AI Literacy Fellowship for African Policymakers.
- International Labour Organization. (2019). *Work for a brighter future: Report of the Global Commission on the Future of Work*. ILO. <https://www.ilo.org/publications/work-brighter-future>
- International Labour Organization & Organisation for Economic Co-operation and Development. (2026). *Compendium of best practices for a human-centered development and use of artificial intelligence in the world of work*. ILO. <https://www.ilo.org/publications/compendium-best-practices-human-centered-development-and-use-artificial>
- Lane, M., & Saint-Martin, A. (2021). *The impact of artificial intelligence on the labour market: What do we know so far?* (OECD Social, Employment and Migration Working Paper No. 256). OECD Publishing. <https://doi.org/10.1787/7c895724-en>
- UNESCO. (2022). *Recommendation on the ethics of artificial intelligence*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>
- United Nations Conference on Trade and Development. (2025). *Technology and innovation report 2025: Inclusive artificial intelligence for development*. United Nations. <https://unctad.org/publication/technology-and-innovation-report-2025>

Editorial note on sources

This is the only chapter in the volume submitted with its own reference list, and the editors have preserved every source the authors chose. Two entries required resolution. The entry “African Union Commission. (2024).” had been split across two list items, with its title stranded in the following entry; it has been rejoined. The OECD entry was cited as “OECD. (2021). Artificial Intelligence and the Future of Work”, which does not correspond to a publication of that exact title; the OECD working paper by Lane and Saint-Martin on AI and the labour market, published in 2021, is the closest match and is cited in its place.

One point is worth recording in the authors’ favour. The ILO compendium was cited with a 2026 date, which the editors initially took to be an error. On verification the date is correct: the compendium was prepared under the G7 Employment Working Group during Canada’s 2025 presidency and issued in January 2026. The entry has been expanded to record it as a joint ILO and OECD publication, which the submitted citation did not note.

Verification of claims

Table 3.13 Editorial verification of claims.

Claim as written	Section	Editorial verification
“Over 60% of workers are engaged in informal employment”	§2	Internal to the case study. Consistent with the rural population share used across other chapters in this volume.
“Youth unemployment exceeds 30%”	§3.1	Internal to the case study, and consistent with Thematic Areas 2, 6 and 8.
“The Data Protection Act (2021)”	§3.3, §6.3	Internal to the case study. The same instrument is referenced in Thematic Areas 2 and 7, indicating a shared scenario parameter rather than a real statute.
“Lumora’s near-collapse” (fintech cyber fraud)	§3.1	Internal to the case study, drawn from the Fellowship case document.
All Year 3 targets (under 48 hours, over 50%, over 40%, over 80%)	§8.2	Authors’ own policy targets for the case study, not empirical findings.
Alignment with the AU Digital Transformation Strategy and Agenda 2063	§2	Externally supported. Both instruments exist and set out the commitments described; citation supplied by the authors and verified by the editors.

Editorial judgements

No author review cycle was available before publication. Where the submitted material contained defects whose resolution could be established with confidence from the surrounding text, the volume editors have resolved them rather than leaving them in the published record. Each such resolution is listed below. No substantive claim, target, or policy position has been altered.

Table 3.14 Editorial judgements applied to this chapter.

Location	As submitted	Editorial resolution
Six headings	Six headings had lost their opening letter: “gital Sovereignty...”, “uman-Centred Ai...”, “ommunity AI Work-and-Learn Hubs”, “bunto-First Agri-Advisory AI”, “buntia AI Commons...”, and “onclusion & Call to Action”.	Restored to Digital Sovereignty, Human-Centred AI, Community AI Work-and-Learn Hubs, Ubuntu-First Agri-Advisory AI, Ubuntia AI Commons, and Conclusion & Call to Action. Each missing letter is unambiguous from the body text that follows.
Throughout	The country was named inconsistently as “Ubantia” in three places and “Ubuntu” in two places where the state, not the philosophy, was meant.	Standardised to “Ubuntia” throughout, including within the quoted country profile. The philosophy of Ubuntu is retained where the authors mean the philosophy.
§6.4	The paragraph beginning “Accountability will be ensured through annual AI and Future of Work implementation reports...” appeared twice, verbatim and consecutively.	The duplicate was removed.
§7, Risk 2	“If AI tools are designed solely for smartphone users in Kifaru City with high-speed 4G, will exclude the female farmer...” — the clause had lost its subject.	Restored to “...they will exclude the female farmer...”.
Figure 3.1	The submitted governance artwork showed the “Ministry of Education” box twice among the supporting ministries, and omitted one ministry named in the text.	Redrawn with the five distinct supporting ministries the chapter names in §6.1 and Table 3.3: Education and Human Capital, Digital Economy and Innovation, Agriculture and Rural Development, Health, and Finance and National Planning.
Figure 3.3	The submitted implementation-pathway artwork duplicated its entire flow band, repeated “Sectoral AI Interventions” twice, and rendered the closing phrase as “Empowered Wollofore”.	Redrawn as a single four-stage pathway, with the closing phrase reading “an inclusive and empowered workforce”.
Section numbering	Numbering was typed by hand and inconsistent: “3. 2.” and “6. 2.” appeared with a space, and some sections carried numbers while others did not.	Numbering rebuilt on Word heading styles and applied consistently from §1 to §9. Section order is unchanged.
Front matter	A “TABLE OF CONTENTS” heading with no contents beneath it, and a stray page number “1” in the body text.	Both removed as drafting artifacts.

Location	As submitted	Editorial resolution
Working group block	Two authors were listed with the same email address, and all five personal addresses were included.	Emails withheld. Publishing personal email addresses in a permanent, publicly indexed DOI record is a privacy exposure; institutional affiliation is retained. The duplication is noted here so the group can correct its own records. The group mentor and group administrator are credited in the volume's front matter rather than on the chapter opening, because several people who held those roles also served the volume as editors or mentors and are named once there.
Authorship of this chapter	The submitted document names five authors. A later OpenSchool Initiative synthesis publication lists eight contributors for this thematic team, adding Zubairu Ibrahim, Kurabachew Desta Tegegne and Said Ali Said Yunus.	Authorship here follows the submitted document, which is the primary source for this chapter and the text these five authors produced. The discrepancy is recorded rather than silently resolved, so that anyone comparing the two publications can see it was considered. Said Ali Said Yunus is credited as an author of Thematic Area 2, where he appears in the submitted document.

Finance, Trade and Inclusion

Responsible Fintech and Digital Inclusion Framework

Trust First. Scale Fast. Govern Responsibly. Include Everyone.

CASE STUDY

The Republic of Ubertia

WORKING GROUP

Responsible Fintech and Digital Inclusion Working Group

Rushmi Christine Matete · Jackfruit Associates Limited, Kenya

Fridah Jebowen Keitany · Chamber Business Academy, Kenya National Chamber of Commerce and Industry, Kenya

Sani Bako Mohammed · Central Bank of Nigeria, Nigeria

Suggested citation: Matete, R. C., Keitany, F. J., & Mohammed, S. B. (2026). Responsible Fintech and Digital Inclusion Framework. In N. G. Abdulhamid (Ed.), *Eight Capstones, One Country* (pp. xx–xx). OpenSchool Initiative. <https://doi.org/10.5281/zenodo.22066741>

1. Executive Summary

1.1 Why This Policy

Ubutia's rapid expansion of mobile money and fintech has widened access to financial services, but trust has not kept pace with growth. Rising fraud, opaque digital lending, weak consumer redress, uneven digital literacy, and fragmented regional rules now pose systemic risks to inclusion, innovation, and competitiveness under the African Continental Free Trade Area (AfCFTA) Agreement.

This policy treats trust as a public asset. It puts citizen protection first, builds fair and reliable markets, and ensures that financial innovation grows without exposing users to harm.

1.2 What the Policy Delivers by 2030

- 50% reduction in per-user digital financial fraud losses
- 75% of adults confidently using regulated digital financial services
- Gender parity in access to digital finance
- 40% reduction in the cost and time of cross-border digital payments and trade
- Ubutia positioned as a regional benchmark for responsible fintech governance

1.3 How It Works

The framework focuses on five integrated policy pillars: consumer protection, anti-fraud regulation, cross-border harmonisation, digital financial literacy, and responsible AI in trade. Each action names a responsible authority, an oversight body, and an indicative budget.

1.4 Why Act Now

Every year of delay compounds consumer harm and erodes trust. Acting now allows Ubutia to scale fintech safely while shaping regional standards rather than importing them.

Trust First. Scale Fast. Govern Responsibly. Include Everyone.

2. Context and Rationale

2.1 Vision and Strategic Framing

To build a safe, inclusive financial system where every citizen of Ubuntuia — regardless of income, gender, geography, or digital literacy — can participate confidently in the digital economy.

Trust will be embedded into the foundations of Ubuntuia’s digital economy through strong consumer protection, decisive anti-fraud regulation, inclusive digital literacy, AfCFTA-aligned regional integration, and responsible AI governance. Trust is treated not as a constraint on innovation, but as a prerequisite for sustainable scale.

2.2 Alignment with National, Regional, and Global Frameworks

- **Sustainable Development Goals.** This framework advances SDGs 1 (No Poverty), 5 (Gender Equality), 8 (Decent Work and Growth), 9 (Industry and Innovation), 10 (Reduced Inequalities), 16 (Strong Institutions), and 17 (Partnerships).
- **AfCFTA and regional integration.** The policy operationalises AfCFTA’s digital trade ambitions through harmonised e-KYC, licensing, consumer protection, interoperable payment corridors, and SME-focused trade facilitation.
- **Donor and development financing alignment.** The framework is designed to unlock concessional and blended financing from the World Bank (IDA), the African Development Bank, IMF-linked reforms, UNDP, and bilateral partners, with clear outcomes, budgets, and monitoring mechanisms.

3. Problem Analysis

Ubuntia's fintech ecosystem has grown rapidly, driven by mobile money, digital credit, and platform-based financial services. However, safeguards have not kept pace with adoption. Key challenges include:

- Rising cyber fraud and digital scams that disproportionately harm low-income and first-time users
- Opaque digital lending practices and weak disclosure standards
- Fragmented and inaccessible consumer redress mechanisms
- Uneven digital financial literacy, particularly in rural areas and among women and youth
- Regulatory fragmentation that raises the cost of cross-border payments and trade

Without intervention, these trends risk eroding public trust, slowing adoption, discouraging investment, and undermining Ubuntia's competitiveness under AfCFTA.

4. Policy Objectives and Strategic Outcomes

The framework sets five measurable objectives to be achieved by 2030.

Table 4.1 Policy objectives and strategic outcomes by 2030.

Objective	Baseline	Target by 2030
Reduce per-user digital financial fraud losses	Not specified by authors	50% reduction
Increase confidently active digital financial users	48% of adults	75% of adults
Achieve gender parity in access to regulated digital financial services	Not specified by authors	Parity
Reduce cross-border digital payment and trade transaction costs	Not specified by authors	40% reduction
Establish Ubutia as a regional leader in responsible fintech and digital trade governance	Not specified by authors	Regional benchmark status

Baselines marked “not specified by authors” were absent from the submitted text and have not been inferred.

5. Core Policy Pillars, Actions, Responsibilities, and Budgets

The framework is delivered through five integrated pillars. Every action names a responsible authority, an oversight body, an indicative budget, and the public value it creates.




Figure 4.1 The five policy pillars, their goals and indicative budgets.

5.1 Pillar A: Consumer Protection

Policy goal. Ensure every digital finance user understands their rights and responsibilities, and can access fast, fair redress.

Table 4.2 Pillar A: Consumer Protection: actions, responsibilities and budgets.

Policy Action	Responsible Authority	Oversight Body	Indicative Budget	Public Value
Enact Fintech Consumer Bill of Rights	Minister of Justice & Minister of Finance	Parliament; Human Rights Commission	USD 1.2M	Legal clarity, user trust
Establish Digital Financial Ombudsman	Minister of Finance	Auditor-General	USD 8–10M (5 yrs)	Fast dispute resolution
Mandate responsible business conduct	Central Bank Governor	Financial Services Authority	USD 0.5M	Informed consumer choice
Publish complaint resolution metrics	Financial Services Authority Chairperson	Parliament Finance Committee	USD 0.3M	Transparency

5.2 Pillar B: Anti-Fraud and Cybersecurity Regulation

Policy goal. Shift fraud risk from citizens to the institutions best positioned to manage it.

Table 4.3 Pillar B: Anti-Fraud and Cybersecurity Regulation: actions, responsibilities and budgets.

Policy Action	Responsible Authority	Oversight Body	Indicative Budget	Public Value
National Real-Time Fraud Intelligence Exchange	Central Bank Governor	National Cybersecurity Council	USD 6M	Early fraud detection
Risk-tiered cybersecurity and capital rules	Central Bank	Parliament	USD 1M	System stability
Shared liability for platform negligence	Minister of Finance	Courts; Financial Services Authority	USD 0.5M	Fair risk allocation
Annual fintech stress tests	Central Bank	Auditor-General	USD 0.8M per year	Resilience

5.3 Pillar C: Cross-Border Fintech Harmonisation

Policy goal. Enable safe, affordable regional payments and digital trade.

Table 4.4 Pillar C: Cross-Border Fintech Harmonisation: actions, responsibilities and budgets.

Policy Action	Responsible Authority	Oversight Body	Indicative Budget	Public Value
Align e-KYC and licensing with AfCFTA	Minister of Trade & Finance	Cabinet Committee on Regional Integration	USD 1.5M	Market access
Mutual recognition of licences	Central Bank Governor	Parliament	USD 0.7M	Lower compliance costs
Cap and simplify remittance fees	Central Bank	Consumer Protection Agency	USD 0.3M	Cheaper transfers
Cross-border payment corridors	Minister of Trade	Auditor-General	USD 4M	SME trade growth

5.4 Pillar D: Digital Financial Literacy

Policy goal. Treat digital financial literacy as public infrastructure.

Table 4.5 Pillar D: Digital Financial Literacy: actions, responsibilities and budgets.

Policy Action	Responsible Authority	Oversight Body	Indicative Budget	Public Value
Integrate digital finance into curricula	Minister of Education	Curriculum Council	USD 5M	Long-term inclusion
“Know Before You Tap” campaign	Minister of ICT	Parliament	USD 3M	Fraud prevention
Train local officials and agents	Minister of Local Government	Auditor-General	USD 2M	Rural protection
Incentivize fintech user education	Central Bank	Financial Services Authority	USD 1M	Shared responsibility

5.5 Pillar E: Responsible AI in Trade Facilitation

Policy goal. Use AI to reduce trade friction while protecting fairness and accountability.

Table 4.6 Pillar E: Responsible AI in Trade Facilitation: actions, responsibilities and budgets.

Policy Action	Responsible Authority	Oversight Body	Indicative Budget	Public Value
Human-in-the-loop AI mandates	Minister of Science & Technology	Data Protection Authority	USD 0.6M	Accountability
AI explainability and bias audits	Data Protection Authority	Parliament	USD 1.2M	Fair access
Local-language datasets	Minister of Higher Education	Auditor-General	USD 3M	Cultural relevance
Ethical AI trade sandbox	Minister of Trade	Cabinet	USD 4M	Faster exports

6. Governance and Implementation Model

6.1 National Execution Timeline




Figure 4.2 National execution timeline across four phases from Year 1 to Year 5 and beyond.

Table 4.7 National execution timeline.

Phase	Years	Political Lead	Core Deliverables
Foundation & Trust Reset	Year 1	Minister of Finance	Ombudsman live; Bill of Rights enacted; fraud exchange pilot
Protection & Literacy Scale-Up	Years 2–3	Central Bank Governor	Cyber rules enforced; 5 million citizens trained
Regional & AI Pilots	Years 3–4	Minister of Trade	Payment corridors; AI sandbox
Optimization & Leadership	Year 5+	Cabinet	50% fraud reduction; regional benchmark

6.2 Institutional Arrangements and Oversight

Table 4.8 Institutional arrangements and oversight.

Role	Institution
Lead Coordinating Authority	Ministry of Finance
Technical Regulator	Central Bank of Ubuntu
Supporting Ministries	Justice; ICT; Trade; Education; Local Government; Science and Technology
Oversight	Parliament; Auditor-General; Data Protection Authority
Coordination Mechanism	Inter-Ministerial Committee on Responsible Digital Finance

7. Risk and Mitigation Plan

Each risk below is drawn from the challenges the authors identify in §3, and each mitigation is drawn from the policy actions they specify in §5. The framework treats trust as a public asset, so the central risk in every case is the erosion of that trust.

Table 4.9 Risks and mitigation measures, mapped to the policy pillars.

Risk	Description	Mitigation	Pillar
Consumer harm and loss of redress	Fragmented and inaccessible redress mechanisms leave users without recourse, while opaque digital lending practices and weak disclosure standards obscure the terms consumers are agreeing to.	Enact a Fintech Consumer Bill of Rights; establish a Digital Financial Ombudsman for fast dispute resolution; mandate responsible business conduct; publish complaint resolution metrics.	A
Fraud and cyber attack	Rising cyber fraud and digital scams disproportionately harm low-income and first-time users, with risk currently borne by citizens rather than institutions.	National Real-Time Fraud Intelligence Exchange; risk-tiered cybersecurity and capital rules; shared liability for platform negligence; annual fintech stress tests.	B
Regulatory fragmentation	Divergent national rules raise the cost of cross-border payments and trade, undermining competitiveness under AfCFTA.	Align e-KYC and licensing with AfCFTA; mutual recognition of licences; cap and simplify remittance fees; establish cross-border payment corridors.	C
Exclusion through uneven digital literacy	Digital financial literacy is uneven, particularly in rural areas and among women and youth, placing the least-protected users at greatest risk.	Integrate digital finance into curricula; run the “Know Before You Tap” campaign; train local officials and agents; incentivize fintech user education.	D
Algorithmic opacity and unfair access	AI deployed in trade facilitation may make consequential decisions without accountability, or reproduce bias that restricts access for some groups.	Human-in-the-loop AI mandates; AI explainability and bias audits; local-language datasets; an ethical AI trade sandbox.	E
Erosion of public trust through delay	The authors state that every year of delay compounds consumer harm and erodes trust, discouraging adoption and investment.	Phased national execution timeline with a named political lead for each phase; annual State of Digital Finance Trust Report tabled in Parliament; quarterly public dashboards.	§6, §8
Weak accountability in delivery	Policy commitments may not translate into implementation without clear ownership and independent scrutiny.	Every policy action names a responsible authority and an oversight body; oversight is distributed across Parliament, the Auditor-General and the Data Protection Authority; independent evaluation involves civil society and academia.	§5, §6.2

8. Monitoring, Evaluation, and Transparency

- **Annual State of Digital Finance Trust Report.** Tabled in Parliament and Cabinet.
- **Quarterly public dashboards.** Covering fraud, complaints, inclusion, and gender parity.
- **Independent evaluation.** Involving civil society and academia.
- **Evidence-based policy adjustments.** Insulated from political cycles.

8.1 Consolidated Indicators

The following indicators are drawn from the objectives in §4 and the transparency commitments above.

Table 4.10 Consolidated monitoring indicators.

Indicator	Target	Reporting
Per-user digital financial fraud losses	50% reduction by 2030	Quarterly public dashboard
Adults confidently using regulated digital financial services	From 48% to 75% by 2030	Quarterly public dashboard
Gender parity in access to digital finance	Parity by 2030	Quarterly public dashboard
Cost and time of cross-border digital payments and trade	40% reduction by 2030	Annual Trust Report
Complaint resolution performance	Published metrics (Pillar A)	Financial Services Authority
Regional standing in responsible fintech governance	Regional benchmark by Year 5+	Annual Trust Report

This table consolidates targets stated by the authors in §1.2, §4 and §5. No new indicators or targets have been introduced.

9. Closing Commitment

This framework is a social contract for the digital age. It protects citizens, empowers innovators, reassures investors, and strengthens democratic institutions.

Trust is not assumed. It is built, funded, governed, measured, and earned.

References

The submitted chapter did not include a reference list. The volume editors have identified authoritative sources for the external frameworks and instruments the chapter invokes. Every URL and DOI below was checked and resolved successfully at the time of publication.

- African Union. (2018). *Agreement establishing the African Continental Free Trade Area*. African Union. <https://au.int/en/treaties/agreement-establishing-african-continental-free-trade-area>
- African Union. (2020). *The digital transformation strategy for Africa (2020–2030)*. African Union. <https://au.int/en/documents/20200518/digital-transformation-strategy-africa-2020-2030>
- Alliance for Financial Inclusion. (2016). *Denarau Action Plan: The AFI network commitment to gender and women's financial inclusion*. AFI. <https://www.afi-global.org/publications/denarau-action-plan-the-afi-network-commitment-to-gender-and-womens-financial-inclusion/>
- Demirgüç-Kunt, A., Klapper, L., Singer, D., & Ansar, S. (2022). *The Global Findex Database 2021: Financial inclusion, digital payments, and resilience in the age of COVID-19*. World Bank. <https://doi.org/10.1596/978-1-4648-1897-4>
- G20/OECD. (2011). *G20 high-level principles on financial consumer protection*. OECD Publishing. <https://www.oecd.org/daf/fin/financial-education/g20highlevelprinciplesonfinancialconsumerprotection.htm>
- United Nations. (2015). *Transforming our world: The 2030 Agenda for Sustainable Development* (A/RES/70/1). United Nations. <https://sdgs.un.org/2030agenda>

Editorial note on sources

No references were submitted with this chapter. The sources above were located by the volume editors to support the external instruments the chapter relies on: the AfCFTA Agreement and the African Union digital transformation strategy invoked in §2.2 and Pillar C; the UN Sustainable Development Goals cited in §2.2; the G20/OECD financial consumer protection principles underpinning Pillar A; the AFI Denarau Action Plan on gender and financial inclusion, which corresponds to the gender parity objective; and the World Bank Global Findex, the standard source for digital financial inclusion baselines. They are offered as supporting apparatus; the authors' analysis and wording are unchanged.

Verification of claims

Table 4.11 Editorial verification of claims.

Claim as written	Section	Editorial verification
“Increase confidently active digital financial users from 48% to 75% of adults”	§4	Internal to the case study. Ubuntu is the Fellowship’s fictional country, so the 48% baseline is a scenario parameter rather than a measured national statistic. No external source is cited.
All indicative budgets (USD 0.3M to USD 10M per action)	§5	Authors’ own estimates for the case study. Presented as indicative budgets rather than sourced costings.
AfCFTA digital trade ambitions; harmonised e-KYC and licensing	§2.2, Pillar C	Externally supported. The AfCFTA Agreement and the AU Digital Transformation Strategy set out these ambitions. Citations supplied by the editors.
SDGs 1, 5, 8, 9, 10, 16 and 17	§2.2	Externally supported. Citation to the 2030 Agenda supplied by the editors.
Named financing partners (World Bank IDA, AfDB, IMF, UNDP)	§2.2	Institutions exist and operate the financing instruments described. The chapter does not claim any commitment has been secured.
Targets for 2030 (50% fraud reduction, gender parity, 40% cost reduction)	§1.2, §4	Authors’ own policy targets for the case study, not empirical findings.

Editorial judgements

No author review cycle was available before publication. Where the submitted text contained defects or structural gaps whose resolution could be established with confidence from the surrounding argument, the volume editors have resolved them rather than leaving them in the published record. Each such resolution is listed below. No substantive claim, target, budget, or policy position has been altered.

Table 4.12 Editorial judgements applied to this chapter.

Location	As submitted	Editorial resolution
§7 Risk and Mitigation Plan	Absent. The Capstone Handbook's standard structure requires a Risk and Mitigation Plan; the chapter did not contain one.	Constructed from the authors' own material. Every risk restates a challenge identified in §3 or a rationale given in §1.4; every mitigation restates a policy action specified in §5 or an oversight mechanism in §6. No new risk, mitigation, or commitment has been introduced, and the source pillar is named for each row.
Pillar A table	The row "Mandate responsible" was cut mid-phrase, with "business conduct" appearing as the opening cell of a separate table alongside four empty cells.	The two fragments were rejoined into a single row reading "Mandate responsible business conduct", and the stray empty cells removed.
Pillar C table	The header row stood alone in one table while all four data rows sat orphaned in another, leaving the columns unlabelled.	Header and data rows reunited into one table.
Pillar E table	"Human-in-the-loop AI mandates"	Word-break corruption repaired to "Human-in-the-loop AI mandates".
Pillars D and E tables	"Auditor-General" and "Auditor-General"	Both repaired to "Auditor-General".
Pillar B table	"USD / 0.8M/year"	Stray separator removed; reads "USD 0.8M per year".
Pillar A table	"USD 8–10M (5 / yrs)"	Stray separator removed; reads "USD 8–10M (5 yrs)".
§8 Monitoring and evaluation	Four transparency commitments were listed without consolidated indicators or targets.	The targets the authors state in §1.2, §4 and §5 have been gathered into a single indicator table. No new indicators or targets were introduced.
§4 Policy objectives	Five objectives were stated as prose targets without a baseline column.	Presented in tabular form. Where the authors state a baseline (48% of adults) it is recorded; the remaining baselines are marked "not specified by authors" rather than inferred.

Agriculture and Food Systems

AI Policy Blueprint for Climate-Smart Agriculture and Food Security

Frugal AI Advisories, Progressive Data Sovereignty and Inclusive Finance for Smallholder Farmers

CASE STUDY

The Republic of Ubutia

WORKING GROUP

Agriculture and Food Systems Working Group

Abdelhafeez Elemam Gasoum · Affiliation not supplied

Abubakar Hassan Inuwa · Foundation for Sustainable Smallholder Solutions (FSSS), Nigeria

Hassan Faruk Maradun · Usmanu Danfodiyo University Sokoto, Nigeria

Idris Muhammad Shamsudeen · Central Bank of Nigeria, Nigeria

Suleiman Usman Yusuf · Affiliation not supplied

Tukur Garba Arab · Voice of Nigeria, Nigeria

Suggested citation: Gasoum, A. E., Inuwa, A. H., Maradun, H. F., Shamsudeen, I. M., Yusuf, S. U., & Arab, T. G. (2026). AI Policy Blueprint for Climate-Smart Agriculture and Food Security. In N. G. Abdulhamid (Ed.), *Eight Capstones, One Country* (pp. xx–xx). OpenSchool Initiative. <https://doi.org/10.5281/zenodo.22066741>

1. Executive Summary

Ubuntia’s agriculture sector employs about 60% of the national workforce and supports food security, rural livelihoods, and export income. The sector is increasingly vulnerable to climate variability, while digital transformation advances unevenly through donor-led pilots and foreign technology platforms. Artificial intelligence (AI) offers practical tools for climate adaptation, boosting productivity, market access, and financial inclusion; however, without intentional governance, AI adoption risks worsening inequality, damaging farmer trust, and trapping the country in extractive technology pathways.

Policy Objective: to deploy climate-smart, inclusive, and trustworthy AI solutions that improve food security, empower vulnerable farmers, especially women, and safeguard national data sovereignty.



Figure 5.1 Policy blueprint at a glance: the six stages of the Policy Architect’s Journey, from grounded assessment through to intended outcomes.

This AI Policy Blueprint follows The Policy Architect’s Journey (Building Africa’s AI Future), starting with a grounded assessment of context and risk, then moving to targeted, capacity-aware policy interventions. Agriculture is intentionally chosen as Ubuntia’s pilot sector for responsible AI governance because of its economic importance, climate vulnerability, and suitability for low-risk, hands-on learning. The Blueprint avoids assuming advanced institutions or a national AI strategy, instead focusing on simple delivery models, gradual data sovereignty, and human-centered oversight.

Four mutually reinforcing policy interventions are proposed: (i) a National Frugal AI Climate Advisory Program; (ii) an Agricultural AI Regulatory Sandbox; (iii) a Progressive Agri-Data Commons Framework; and (iv) an Inclusive AI Finance Program for Women Farmers. These interventions are sequenced across short-, medium-, and long-term horizons, anchored in existing institutions, and supported by clear monitoring and evaluation mechanisms.

2. Context and Rationale

Agriculture remains the backbone of Ubuntia’s economy, employing most of the workforce and supporting the rural population within a total of 42 million residents, with a median age of 20 years. The sector is mostly composed of smallholder farmers who rely on rain-fed agriculture and are therefore highly susceptible to climate shocks, such as unpredictable rainfall, droughts, and flooding. Recent seasons have shown that climate variability directly leads to decreases in crop yields, income fluctuations, and heightened food insecurity.




Figure 5.2 Ubuntia agriculture snapshot: the six structural parameters that frame this blueprint.

Ubuntia’s digital landscape amplifies these risks. Rural internet access remains below 50%, smartphone and formal financial service access is uneven, and digital literacy gaps, especially for women, remain. Despite these limitations, AI tools are already being adopted in the agricultural sector through donor-funded projects, multinational tech providers, and startup pilots. These efforts often depend on foreign cloud infrastructure, imported models trained on non-local data, and interfaces that exclude indigenous languages.

These conditions are not unique to Ubuntia. The Intergovernmental Panel on Climate Change (2022) projects declining yields for staple cereals across sub-Saharan Africa under all warming scenarios, with rain-fed smallholder systems identified as the most exposed livelihood group on the continent. The landscape of digital agricultural services is likewise well documented: a joint survey by the Technical Centre for Agricultural and Rural Cooperation and GSMA (2019) mapped more than 600 digital agriculture services operating across the region and found that SMS- and voice-based advisory channels dominate precisely because they work on low-end handsets and do not presume literacy.

3. Problem Analysis

The main policy challenge is not whether AI will be used in agriculture but how it should be governed. Without clear policy guidance, AI adoption could become chaotic, exploitative, and exclusive. On the other hand, premature or overly strict regulations might hinder the development of key climate adaptation tools. This Blueprint recommends a phased, capacity-aware governance approach that balances innovation with protections for the public interest.



Figure 5.3 AI risk–opportunity matrix: outcomes for Ubuntu agriculture plotted against AI adoption and governance strength.

3.1 Key Risks

Data Sovereignty and Dependency. With a large share of cloud and AI services owned by foreign entities, farmer-generated data, such as soil conditions, yields, and weather patterns, risks being collected, processed overseas, and monetized without providing proper national benefits. This reliance could threaten food security if these services are pulled back or their prices are increased.

Algorithmic Bias and Contextual Mismatch. Imported AI models trained on crops and farming systems from the Global North may provide inaccurate or harmful advice when applied to Ubuntu's indigenous crops and local ecological conditions. Limited support for the local language further marginalizes rural farmers.

Automated Exclusion. AI tools that require smartphones, formal bank accounts, or land ownership records risk excluding women and smallholders in informal economies, reinforcing existing inequalities.

Trust and Liability Failures. Incorrect AI-generated climate or planting advice could trigger widespread crop losses. In a context of fragile institutional trust, such failures could lead to the rejection of digital tools entirely.

The second and third of these risks are well attested in the literature. Birhane (2020) documents how AI systems built on Global North assumptions and training data are structurally mismatched to African social and ecological realities, while Nekoto et al. (2020)

show that the overwhelming majority of African languages have no machine translation resources at all — a direct constraint on any advisory service that intends to reach farmers in the language they actually speak.

3.2 Strategic Opportunities

Climate Adaptation through Frugal AI. AI-driven climate intelligence can significantly improve planting decisions when delivered through low-tech channels such as SMS and voice.

Inclusive Finance. Machine learning using alternative data, such as mobile money histories, can expand access to credit for farmers lacking formal collateral, particularly women.

Market Transparency. AI-enabled logistics and pricing systems can reduce post-harvest losses and strengthen farmer bargaining power.

Sovereign Innovation. Developing AI tools in local languages and contexts positions Ubuntia as a producer, not merely a consumer, of agricultural AI solutions.

4. Policy Objectives

4.1 Policy Objective

Deploy climate-smart, inclusive, and trustworthy AI in agriculture to strengthen food security, improve farmer resilience, and protect national data sovereignty.

4.2 Guiding Principles

Table 5.1 The five guiding principles governing AI deployment in Ubuntu agriculture.

Principle	What it requires
Food Security First	AI systems that influence agricultural decisions are considered critical infrastructure.
Inclusion by Design	Tools must intentionally support women, rural farmers, and users with low literacy.
Progressive Sovereignty	Data governance develops gradually through practical, phased steps rather than immediate, strict mandates.
Human Oversight	AI supports instead of replacing extension services and farmer judgment.
Innovation with Guardrails	Experimentation occurs through sandboxes and pilots within defined safeguards.

The third principle is the one most often skipped. Canfield and Ntambirweki (2024) argue that agricultural data governance in African contexts has tended to move directly from open-data enthusiasm to platform capture, without an intervening stage in which farmers hold enforceable rights over the data they generate. A phased approach is what creates room for that stage to exist.

5. Proposed Interventions




Figure 5.4 The four policy interventions, with lead institution and delivery timeline.

5.1 National Frugal AI Climate Advisory Program

This program will provide AI-supported weather forecasts, planting guidance, and early-warning alerts via SMS and voice services in local languages. Centralized data processing will be combined with validation by agricultural extension officers before distribution.

Lead Institution: Ministry of Agriculture and Rural Development. Timeline: pilot (6–12 months); nationwide scale-up (1–2 years).

5.2 Agricultural AI Regulatory Sandbox

A sector-specific sandbox will enable startups, donors, and public agencies to test AI tools under regulatory oversight. Participation will require adherence to an Agricultural AI Ethics Protocol that addresses data privacy, inclusion, transparency, and human oversight.

Lead Institution: Ministry of ICT and Digital Economy. Timeline: framework design (6 months); ongoing pilot cohorts.

The sandbox model has an established evidence base in exactly this kind of setting. The Consultative Group to Assist the Poor (2017) found that time-limited, supervised live testing is particularly suited to regulators facing capacity gaps and fast-moving technology, while cautioning that a sandbox is not a substitute for a settled regulatory framework.

5.3 Progressive Agri-Data Commons Framework

This framework will gradually implement the localization and management of agricultural data, starting with public datasets and progressing toward a shared Agri-Data Commons accessible to approved users.

Lead Institution: Data Protection Authority. Timeline: standards and localization (1–2 years); full commons (3–5 years).

5.4 Inclusive AI Finance for Women Farmers

AI-driven credit scoring using alternative data will be piloted to expand access to finance for women farmers, with strict transparency and grievance mechanisms.

Lead Institution: Central Bank of Ubuntia. Timeline: pilot (12 months); expansion (2–3 years).

The transparency and grievance requirements are not incidental to this intervention; they are what distinguishes it from the digital credit products already circulating in the region. Kaffenberger, Totolo and Soursourian (2018) found that roughly half of digital credit borrowers surveyed in Kenya and Tanzania repaid late, that a substantial minority did not understand the cost of the loans they had taken, and that borrowers skewed young, urban and male. The gender dimension is equally documented: the Global Findex database records a persistent gap of around twelve percentage points between men’s and women’s account ownership in sub-Saharan Africa (Demirgüç-Kunt et al., 2022), and the FAO (2024) identifies women’s restricted access to credit and digital services as both a productivity gap and a climate-vulnerability multiplier.

5.5 Policy Intervention Summary

Table 5.2 Policy intervention summary, in the Capstone Handbook's prescribed format.

Intervention	Description	Implementing Body	Resources Needed	Timeline	Indicators
1. National Frugal AI Climate Advisory Program	AI-supported weather forecasts, planting guidance and early-warning alerts delivered via SMS and voice in local languages, with centralized data processing validated by agricultural extension officers before distribution.	Lead: Ministry of Agriculture and Rural Development	SMS and voice delivery channels; local-language advisory content; centralized data processing; extension-officer validation capacity †	Pilot 6–12 months; nationwide scale-up 1–2 years	• 500,000+ farmers receiving reliable advisories ‡ • 20–25% reduction in climate-related crop losses ‡
2. Agricultural AI Regulatory Sandbox	A sector-specific sandbox enabling startups, donors and public agencies to test AI tools under regulatory oversight, conditional on adherence to an Agricultural AI Ethics Protocol covering data privacy, inclusion, transparency and human oversight.	Lead: Ministry of ICT and Digital Economy	Agricultural AI Ethics Protocol; supervisory and cohort administration capacity †	Framework design 6 months; ongoing pilot cohorts	• Increased farmer trust in AI-enabled services ‡
3. Progressive Agri-Data Commons Framework	Phased localization and management of agricultural data, beginning with public datasets and progressing toward a shared Agri-Data Commons accessible to approved users.	Lead: Data Protection Authority	Data localization standards; stewardship arrangements; approved-user access controls †	Standards and localization 1–2 years; full commons 3–5 years	• Data governance compliance ‡
4. Inclusive AI Finance for Women Farmers	AI-driven credit scoring using alternative data, piloted to expand access to finance for women farmers, with strict transparency and grievance mechanisms.	Lead: Central Bank of Ubutia	Alternative data sources such as mobile money histories; transparency and grievance mechanisms †	Pilot 12 months; expansion 2–3 years	• 25,000 women farmers accessing AI finance ‡ • Inclusion metrics ‡

† Resources marked with a dagger are drawn from the authors' own descriptions of each intervention in §5.1–5.4; the submission does not contain a separate resourcing schedule. ‡ Indicators marked with a double dagger are the chapter's own success metrics (§6.3) and monitoring priorities (§8), mapped by the editors to the intervention each most directly measures. Neither has been invented.

6. Governance and Implementation Model

6.1 Governance and Stakeholder Framework

Implementation will rely on existing institutions coordinated by an Interim AI-in-Agriculture Working Group led by the Ministry of Agriculture. Agricultural extension services will act as trusted intermediaries between AI systems and farmers. Donors and private actors will operate within nationally established guidelines, while universities and startups will support innovation and capacity building.




Figure 5.5 Governance and stakeholder framework, showing the coordinating body, lead institutions, delivery actors and intended beneficiaries.

6.2 Why Act Now

Climate shocks are already eroding livelihoods. AI adoption is accelerating without governance. Delay locks in dependency and exclusion.

6.3 What Success Looks Like

Table 5.3 The chapter's stated success metrics.

Measure	Target
Reduction in climate-related crop losses	20–25%
Farmers receiving reliable advisories	500,000+
Women farmers accessing AI finance	25,000
Farmer trust in AI-enabled services	Increased
Overall national position	Transition from “AI testing ground” to regional leader

6.4 Implementation Roadmap and Sequencing




Figure 5.6 Phased implementation roadmap across short-, medium- and long-term horizons.

Short Term (0–12 months). Pilot frugal AI advisories; establish a sandbox; form an interim working group.

Medium Term (1–3 years). Scale successful pilots, implement data governance standards, and expand inclusive finance.

Long Term (3–5 years). Institutionalize the Agri-Data Commons and incorporate lessons into a national AI strategy.

This sequencing is consistent with the continental policy direction. The African Union’s Continental Artificial Intelligence Strategy, endorsed by the Executive Council in July 2024, calls for Africa-centric, development-focused national approaches to AI rather than the wholesale importation of external regulatory models, and sits within the broader Digital Transformation Strategy for Africa (2020–2030).

7. Risk and Mitigation Plan

The submission analysed risk thoroughly in its opening sections but did not pair each risk with a mitigation. The table below closes that gap. Every mitigation is drawn from the chapter's own guiding principles (§4.2) and policy interventions (§5); none has been introduced from outside the submitted text.

Table 5.4 Risk and mitigation plan, constructed by the editors from the chapter's own principles and interventions.

Risk	How it could materialise	Mitigation drawn from this chapter	Lead
Data sovereignty and dependency	Farmer-generated data on soils, yields and weather is collected, processed offshore and monetized without national benefit; food security is exposed if foreign services withdraw or raise prices.	The Progressive Agri-Data Commons Framework (§5.3) phases in localization and stewardship of agricultural data, beginning with public datasets. The Progressive Sovereignty principle (§4.2) commits to practical, staged steps rather than mandates the state cannot yet enforce.	Data Protection Authority
Algorithmic bias and contextual mismatch	Models trained on Global North crops and farming systems give inaccurate or harmful advice for indigenous crops; limited local language support marginalizes rural farmers.	Advisories are issued in local languages and validated by extension officers before distribution (§5.1). The Agricultural AI Ethics Protocol governing sandbox participation (§5.2) requires transparency and human oversight, and the sandbox itself provides a supervised setting in which contextual failure can surface before national deployment.	Ministry of Agriculture; Ministry of ICT and Digital Economy
Automated exclusion	Tools requiring smartphones, formal bank accounts or land title exclude women and smallholders in informal economies, reinforcing existing inequalities.	Delivery is by SMS and voice rather than smartphone (§5.1). Credit scoring uses alternative data such as mobile money histories in place of formal collateral, and is targeted at women farmers (§5.4). Inclusion by Design (§4.2) requires tools to support women, rural farmers and users with low literacy.	Central Bank of Ubuntu; Ministry of Agriculture
Trust and liability failures	Incorrect AI-generated climate or planting advice triggers widespread crop losses; in a context of fragile institutional trust, this leads to wholesale rejection of digital tools.	Human Oversight (§4.2) makes AI a support to, not a replacement for, extension services and farmer judgment, and extension-officer validation stands between the model and the farmer (§5.1). The finance pilot carries grievance mechanisms (§5.4), and annual public reporting (§8) exposes performance to scrutiny.	Ministry of Agriculture; extension services

8. Monitoring and Evaluation Framework

Monitoring will emphasize outcomes rather than technology adoption, concentrating on yield consistency, inclusion indicators, data governance compliance, and farmer trust. Annual public reports will promote transparency and facilitate adaptive learning.




Figure 5.7 Monitoring, evaluation and learning framework, structured as a results chain from inputs to impact.

The framework deliberately measures the results of AI deployment rather than its extent. Counting deployed tools or registered users would show activity; measuring yield stability, inclusion, compliance and trust shows whether the policy objective in §4.1 is actually being met. The annual public report is the mechanism that makes this visible outside government.

9. Conclusion and Call to Action

This AI Policy Blueprint explains how Ubuntia can govern AI responsibly and practically by focusing on agriculture, where risks and benefits are most immediate. By adopting affordable strategies, creating forward-thinking data policies, and promoting inclusive innovation within existing institutions, the Blueprint presents AI as a tool for resilience rather than disruption. The lessons from agricultural uses can help shape a future national AI strategy grounded in trust, inclusion, and sovereignty.

9.1 Decision Implications for Policymakers

- Endorse agriculture as Ubuntia's pilot sector for AI governance and learning.
- Approve time-limited pilots and regulatory sandboxes instead of comprehensive AI legislation.
- Mandate inclusion, human oversight, and data stewardship as essential safeguards.
- Use early results to shape a broader national AI strategy across sectors.

References

This chapter was submitted without a reference list. The sources below were located by the editors to support the claims the authors make, and every locator was checked before inclusion. They are offered as the evidence base the chapter draws on implicitly, not as citations the authors themselves supplied.

- African Union. (2020). *The digital transformation strategy for Africa (2020–2030)*. African Union Commission. <https://au.int/en/documents/20200518/digital-transformation-strategy-africa-2020-2030>
- African Union. (2024). *Continental artificial intelligence strategy*. African Union Commission. <https://au.int/en/documents/20240809/continental-artificial-intelligence-strategy>
- Birhane, A. (2020). Algorithmic colonization of Africa. *SCRIPTed: A Journal of Law, Technology & Society*, 17(2), 389–409. <https://doi.org/10.2966/scrip.170220.389>
- Canfield, M., & Ntambirweki, B. (2024). Datafying African agriculture: From data governance to farmers' rights. *Development*, 67(1–2), 5–13. <https://doi.org/10.1057/s41301-024-00405-7>
- Consultative Group to Assist the Poor. (2017). *Regulatory sandboxes and financial inclusion* (Working paper). CGAP. <https://www.cgap.org/research/publication/regulatory-sandboxes-and-financial-inclusion>
- Demirgüç-Kunt, A., Klapper, L., Singer, D., & Ansar, S. (2022). *The Global Findex database 2021: Financial inclusion, digital payments, and resilience in the age of COVID-19*. World Bank. <https://www.worldbank.org/en/publication/globalindex>
- Food and Agriculture Organization of the United Nations. (2024). *The status of women in agrifood systems in sub-Saharan Africa: Overview*. FAO. <https://openknowledge.fao.org/handle/20.500.14283/cc5343en>
- Intergovernmental Panel on Climate Change. (2022). Africa. In H.-O. Pörtner, D. C. Roberts, M. Tignor, E. S. Poloczanska, K. Mintenbeck, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, & B. Rama (Eds.), *Climate change 2022: Impacts, adaptation and vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (pp. 1285–1456). Cambridge University Press. <https://doi.org/10.1017/9781009325844.011>
- Kaffenberger, M., Totolo, E., & Soursourian, M. (2018). *A digital credit revolution: Insights from borrowers in Kenya and Tanzania* (Working paper). CGAP. <https://www.cgap.org/research/publication/digital-credit-revolution-insights-borrowers-in-kenya-and-tanzania>
- Nekoto, W., Marivate, V., Matsila, T., Fasubaa, T., Fagbohunge, T., Akinola, S. O., Muhammad, S., Kabenamualu, S. K., Osei, S., Sackey, F., Niyongabo, R. A., Macharm, R., Ogayo, P., Ahia, O., Berhe, M. M., Adeyemi, M., Mokgesi-Seling, M., Okegbemi, L., Martinus, L., ... Bashir, A. (2020). Participatory research for low-resourced machine translation: A case study in African languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020* (pp. 2144–2160). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.195>
- Smart Africa. (2021). *Blueprint: Artificial intelligence for Africa*. Smart Africa Secretariat. <https://smartafrica.org/project-post/artificial-intelligence-for-africa/>
- Technical Centre for Agricultural and Rural Cooperation, & GSMA. (2019). *The digitalisation of African agriculture report 2018–2019*. CTA. <https://cgspace.cgiar.org/items/fb60e627-208f-4ae1-aba1-40bc2054e856>

Editorial note on sources

The submission contained no citations and no reference list. Rather than publish a policy chapter with no evidence base, the editors located sources for each substantive external claim the chapter makes and added in-text references at the points where the chapter relies on them. Every URL and DOI above was opened and confirmed to resolve to the work cited; the Canfield and Ntambirweki article was confirmed through the Crossref metadata registry after the publisher's site returned an automated-access refusal rather than a genuine error.

Two candidate sources were rejected during this process. A GSMA AgriTech landscape report was dropped in favour of the CGSpace deposit of the CTA/GSMA study, because GSMA's own copies of the relevant PDFs no longer resolve. A separately published African Union Continental AI Strategy could not initially be confirmed by the research assistance used for this chapter; the editors subsequently located it on the African Union's own document repository and verified the July 2024 Executive Council endorsement, so it is cited here.

Verification of claims

Table 5.5 Editorial verification of claims.

Claim as written	Section	Editorial verification
"Agriculture employs about 60% of the national workforce"; "42 million residents"; "median age of 20 years"	§1, §2	Internal to the case study. These are shared scenario parameters and are consistent with the figures used in Thematic Areas 1, 2, 3 and 6.
"Rural internet access remains below 50%"	§2	Internal to the case study. Other chapters give national internet penetration as 48%, which is consistent with a lower rural figure.
Climate variability reducing crop yields and raising food insecurity	§2	Externally supported. IPCC (2022) projects declining staple cereal yields across sub-Saharan Africa and identifies rain-fed smallholder systems as the most exposed.
Imported models mismatched to indigenous crops; limited local language support	§3.1	Externally supported. Birhane (2020) on structural mismatch; Nekoto et al. (2020) on the near-absence of African language resources in machine translation.
SMS and voice as effective low-tech advisory channels	§3.2, §5.1	Externally supported. CTA and GSMA (2019) mapped over 600 digital agriculture services in the region and found SMS and voice channels predominate.
Alternative-data credit scoring expanding access for women farmers	§3.2, §5.4	Externally supported, with qualification. Demirgüç-Kunt et al. (2022) document the gender gap in account ownership and FAO (2024) the credit gap for women in agrifood systems; Kaffenberger et al. (2018) document real risks of late repayment and poor cost transparency in digital credit, which is why the chapter's transparency and grievance safeguards matter.
"20–25% reduction in climate-related crop losses"; "500,000+ farmers receiving reliable advisories"; "25,000 women farmers accessing AI finance"	§6.3	These are stated as policy targets, not findings, and are presented as such. No external source is cited or required.
The Ministry of ICT and Digital Economy, Data Protection Authority and Central Bank of Ubutia	§5, §6.1	Internal to the case study. The same institutions appear in Thematic Areas 1, 2 and 4, indicating shared scenario architecture.

Editorial judgements

This volume is being published without a further round of author review. Where the submission contained defects, the editors have resolved them on the evidence of the chapter's own text rather than leaving queries for authors who will not see them. Every intervention is recorded below.

Table 5.6 Editorial judgements applied to this chapter.

Location	As submitted	Editorial resolution
Chapter structure	Ordered as Executive Summary, Context and Problem Definition, Risk and Opportunity Analysis, Policy Objective and Guiding Principles, Policy Interventions, Governance, Implementation Roadmap, Monitoring, Conclusion, Decision Implications.	Resequenced to the Capstone Project Handbook's nine-section blueprint used throughout this volume. The risk analysis now supplies §3 and §7, the roadmap joins governance in §6, and the decision implications close §9. No content was dropped.
Risk section	Four risks were identified and described, but no mitigation was paired with any of them, and the Handbook structure requires a risk and mitigation plan.	A mitigation plan was constructed in §7 by mapping each stated risk onto the guiding principles and interventions the authors themselves propose. Nothing was introduced from outside the submission, and the construction is declared in the table's caption and preamble.
References	No citations and no reference list.	Twelve sources located by the editors and verified, with in-text references added at the points where the chapter relies on them. Documented in the editorial note above.
All seven figures	Supplied at approximately 680 pixels wide — about 108 dpi at page width — in greyscale, and inconsistent with the colour house style used across the volume. The monitoring framework's lowest band overlapped the tier above it.	All seven redrawn at print resolution in the volume's house style. Content is the authors' own; this was the cleanest figure set in the volume and no factual error was found in any of them.
Figure 5.1	A six-column horizontal strip. At A4 page width the body text would have set at roughly 4.5 point.	Rotated to a vertical six-stage spine. The sequence and all content are unchanged; the rotation buys the line length needed to set the text legibly.
Figure 5.5, governance map	A hub-and-spoke diagram centred on the Ministry of Agriculture, with the Interim Working Group drawn as one spoke among peers. Farmers did not appear at all.	Redrawn as four tiers. The chapter's own §6.1 states that the Working Group coordinates implementation and is led by the Ministry, so it is placed above rather than beside the institutions it coordinates. Farmers are added as the intended beneficiaries, reached through extension services — the role the same paragraph assigns them. Every actor the authors named is retained.
Figure 1 caption	"Figure 1Policy Blueprint at a Glance", with no space after the figure number.	Corrected, and all captions set in the volume's house format.

Location	As submitted	Editorial resolution
Working group block	Two of the six authors — Abdelhafeez Elemam Gasoum and Suleiman Usman Yusuf — had the template’s placeholder text “Institutional/Organisational Affiliation”, “Country” and “Email Address” left in place of their details. The other four supplied personal email addresses.	The placeholder text is not published as though it were an affiliation; those two entries read “Affiliation not supplied”. The editors had no basis on which to infer an institution and did not guess. All emails withheld: publishing personal addresses in a permanent, publicly indexed DOI record is a privacy exposure, and institutional affiliation is retained where given.
§3.1, §3.2 lists	Numbered as “i.”, “ii.”, “iii.”, “iv.” with each marker stranded on its own line, and several items severed mid-sentence — artifacts of the conversion from the compendium.	Reflowed as continuous prose with the lead term set in the running text. Wording is unchanged.
§4.2, guiding principles	A numbered list broken across paragraph boundaries mid-item, so that “AI systems that influence agricultural decisions are considered” and “critical infrastructure” fell in separate paragraphs.	Rejoined and set as a table. Wording is unchanged.
§6, “WHY ACT NOW & SUCCESS METRICS”	A slide-style heading block with bulleted fragments beneath it.	Recast as §6.2 Why Act Now and §6.3 What Success Looks Like, with the metrics set as a table. Content is unchanged.
Throughout	A page-footer artifact reading “PUBLIC” appeared on every page, together with bare page numbers carried over from the source document.	Removed. Pagination is now handled by the volume’s own footer.

Climate Resilience and Urban Systems

AI Policy Blueprint for Climate-Resilient, Data-Driven Urban Systems

An Africa-Driven Governance Framework for Smart Cities, Clean Energy, and Climate Finance

CASE STUDY

The Republic of Ubuntia

WORKING GROUP

Climate Resilience and Urban Systems Working Group

Ismail Musa Umar · Nigerian Television Authority, Nigeria

Maryam Rabi Aliyu · Cyprus International University, Mersin, Türkiye

Suggested citation: Umar, I. M., & Aliyu, M. R. (2026). AI Policy Blueprint for Climate-Resilient, Data-Driven Urban Systems. In N. G. Abdulhamid (Ed.), *Eight Capstones, One Country* (pp. xx–xx). OpenSchool Initiative. <https://doi.org/10.5281/zenodo.22066741>

1. Executive Summary

Artificial Intelligence is diffusing across the globe faster than any prior general-purpose technology. Yet diffusion alone does not equal development. Without deliberate governance, Africa risks becoming a passive testing ground for AI systems designed elsewhere, trained on foreign data, optimized for external markets, and misaligned with local realities.

This policy blueprint proposes a ground-up, Africa-authored AI governance framework, anchored in climate resilience and urban systems. Using the Republic of Ubuntia as a representative case, it demonstrates how AI can be deployed responsibly across smart transport, renewable energy, buildings, vertical farming, waste systems, and carbon markets, provided coordination precedes implementation.

The blueprint integrates:

- Smart city infrastructure
- Renewable energy systems
- AI-enabled carbon credit generation
- Institutional coordination mechanisms
- Risk-based AI governance
- Climate finance reinvestment loops

At its core lies a guiding principle:

“If I don’t understand it, I will not adopt it.”

This document prioritizes clarity, accountability, and implementation, ensuring AI policy is not only visionary, but executable.

2. Context and Rationale

2.1 Introduction: Why This Policy, Why Now

Africa stands at a paradoxical moment. On one hand, the continent hosts the world's youngest population, rapidly urbanizing cities, and growing digital adoption. On the other, it faces infrastructure deficits, climate vulnerability, and institutional fragmentation.

AI enters this landscape not as a neutral tool, but as a force amplifier. Without a national AI strategy and risk-based oversight, Ubuntia risks:

- Importing opaque AI systems trained on Western datasets
- Embedding cultural bias into public services
- Deepening urban inequality
- Missing climate finance opportunities tied to data credibility

Conversely, with intentional design, AI can:

- Optimize urban systems
- Accelerate renewable energy integration
- Generate verifiable carbon credits
- Mobilize climate finance
- Create dignified work for youth

The difference lies not in technology, but in governance.

2.2 Country Context: Ubuntia as Africa in Miniature

Ubuntia is a mid-sized Sub-Saharan nation of 42 million people, with over 60% living in rural areas, yet urbanization is accelerating. Secondary cities are absorbing migrants faster than planning systems can respond, spanning coastal plains and mountainous interiors. Its capital, Lumora, anchors government and finance, while Kifarú City has emerged as a technology hub driven by youth entrepreneurship.



Figure 6.1 The Republic of Ubuntia, showing the capital Lumora and the technology hub Kifaru City.

Key characteristics:

- Median age: approximately 20 years
- Rapid urbanization
- Agriculture still dominant
- Democratic institutions under strain
- Growing digital services, but uneven infrastructure

Urban systems are under pressure:

- Transport congestion
- Energy demand outpacing supply
- Informal settlements expanding
- Waste management failures
- Rising emissions trajectory

Illustrative baseline indicators:

- Urban transport emissions growing at 8–12% annually
- Renewable energy share of urban electricity: below 30%
- Youth unemployment: over 30%, higher for young women




Figure 6.2 Urban emissions by sector (illustrative). Transport 35%, energy 30%, buildings 20%, waste 15%.

Despite growing data availability (smart meters, satellite imagery, mobile platforms), policies fail to translate evidence into action due to:

- Institutional silos
- Weak municipal capacity
- Lack of coordination frameworks
- AI systems introduced without governance clarity

AI policy in Ubutia must therefore solve real problems, not theoretical ones.

2.3 Why This Matters

AI diffusion mirrors historical inequality. According to global assessments, countries with reliable electricity, broadband, and skills dominate frontier AI development, while others risk dependency.

Urban systems will drive over 70% of future emissions growth in rapidly urbanizing African economies. Cities are both the problem and the solution. If AI is embedded into transport optimization, energy demand forecasting, building efficiency, and waste sorting and recovery, then cities can become engines of climate mitigation rather than emissions. But only if data is clean, institutions are coordinated, and citizens trust the system.

For Ubutia, the danger is not technological lag. It is policy misalignment.

Table 6.1 *Ubuntu's choice: the risk of policy misalignment against the opportunity of Africa-driven AI leadership.*

The Risk: Policy Misalignment	The Opportunity: Africa-Driven Leadership
If AI systems are procured without understanding	Design systems around local realities
If AI systems use foreign-trained models without localization	Retain data sovereignty
If AI systems operate without public oversight	Build institutions rather than bypass them
Then trust collapses, adoption fails, and public value is lost	Turn climate responsibility into economic opportunity

3. Problem Analysis

3.1 Fragmented Governance

Across Africa, over 60% of digital public projects fail to scale due to weak inter-ministerial coordination rather than lack of technology.

3.2 Imported Intelligence

Most AI systems deployed are:

- Trained on non-African data
- Blind to local dialects, tonation, and norms
- Optimized for contexts with stable electricity and broadband

Bias is not accidental — it is structural.

3.3 Climate Finance Gap

Africa receives less than 5% of global climate finance, despite bearing disproportionate climate risk. One key barrier is the lack of trusted Measurement, Reporting, and Verification (MRV) systems. AI could solve this, if governed correctly.

4. Policy Objectives

4.1 Vision

To build Africa-driven AI systems that power climate-resilient urban development while preserving sovereignty, trust, and inclusion.

4.2 Core Principles

1. **Coordination before implementation** — No AI deployment without institutional alignment.
2. **Actionable policy outputs** — Every recommendation must specify who does what, when, and with what resources.
3. **Africa-driven leadership** — Models, data, and governance reflect African context, not imported assumptions.
4. **Data sovereignty and trust** — Data is publicly owned, transparently governed, and ethically used.

5. **Risk-based AI governance** — High-risk applications face higher scrutiny and safeguards.
6. **Measurable impact and reinvestment** — No impact claimed without verification. No finance without measurement.

In summary: coordination before implementation; public ownership of strategic data; risk-based AI deployment; explainability over complexity; local relevance over imported intelligence; and finance that follows verified impact.

5. Proposed Interventions

5.1 Smart City Infrastructure and AI Applications

Four urban domains anchor the blueprint. Each generates verifiable emissions data that feeds the national MRV pipeline.

Smart Transport

- AI traffic optimization
- Public transit demand prediction
- Autonomous vehicle pilots (controlled zones)

Renewable Energy

- AI-driven grid balancing
- Solar and wind forecasting
- Smart meters for consumption data

Buildings

- Energy efficiency modeling
- Retrofitting prioritization
- Emissions baseline tracking
- Smart homes
- Vertical farming

Waste Systems

- AI-assisted sorting
- Methane capture measurement
- Circular economy analytics

Each system feeds verifiable emissions data into national platforms.

5.2 Data Collection Layer

Table 6.2 Data collection layer: sources and governance requirements.

Data sources	Data governance requirements
IoT sensors	Publicly governed
Smart meters	Privacy-preserving
Satellite imagery	Standardized for MRV
Municipal data systems	—

Clean data is necessary, but insufficient without coordination.

5.3 AI Analytics and MRV

AI functions in this layer include:

- Emissions measurement
- Impact verification
- Fraud detection
- Forecasting future reductions

This layer transforms raw data into trusted climate assets.

5.4 Carbon Credits and Climate Finance

Carbon markets reward verified emissions reductions. AI-enabled MRV reduces verification costs, fraud risk, and time delays.

Table 6.3 Illustrative benchmarks: traditional against AI-enabled MRV.

Metric	Traditional MRV	AI-Enabled MRV
Cost	High, manual	Low, automated
Frequency	Infrequent (e.g., annual)	Continuous, real-time
Scalability	Low	High
Auditing	Paper-based	Digital, auditable ledger

The value chain runs as follows:

7. Urban action across transport, energy, buildings and waste generates emissions reductions.
8. A coordination layer aligns ministries, cities, regulators and auditors before analysis begins.
9. AI-enabled MRV verifies those reductions.
10. The National Carbon Registry issues credits against verified reductions.
11. Climate finance is unlocked and reinvested into resilience.

5.5 National Carbon Registry

A publicly governed carbon registry must be transparent, auditable, and aligned with international standards. AI strengthens credibility, unlocking climate finance.

5.6 Pilot Programmes

Pilot A: Energy Demand Management (Municipal Building Retrofit)

Scope. One municipal building (e.g., city hall or public library) retrofitted with sensors, smart meters, and an AI controller to optimize HVAC, lighting, and storage use.

Objectives:

- Reduce building energy consumption and peak demand.
- Demonstrate occupant-aware control that preserves comfort.
- Produce verifiable kWh and tCO₂e savings for outcome-based contracting.

Core activities:

- Install occupancy sensors, smart thermostats, sub-metering, and a local edge controller.
- Integrate rooftop solar forecast feed and battery dispatch logic.
- Train and deploy an AI model to predict occupancy and optimize HVAC schedules.
- Run 6 months baseline, 12 months active operation, continuous MRV.

Key performance indicators:

- Target 15–25% reduction in annual energy use.
- Peak demand reduction 10–20%.
- Verified savings in kWh and tCO₂e reported monthly.
- Occupant comfort maintained (survey score at or above baseline).

Pilot B: Transport Corridor AI Traffic Optimization

Scope. AI signal optimization on a 2–4 km urban corridor with mixed traffic to reduce idling, emissions, and travel time.

Objectives:

- Reduce vehicle idling and emissions.
- Improve average travel speed and public transit reliability.
- Demonstrate measurable congestion minutes avoided for outcome payments.

Core activities:

- Deploy vehicle and pedestrian detection sensors (camera or radar with privacy filters) or use existing loop detectors.
- Integrate traffic signal controllers with an edge AI optimizer.
- Run A/B testing: baseline weeks against optimized weeks.
- Public dashboard showing congestion and emissions metrics.

Key performance indicators:

- Idling time reduction of at least 20% on the corridor.
- Average travel time reduction of at least 10%.
- Reduction in estimated vehicle emissions (gCO₂e/km) measured by modeled emissions.
- Public transit on-time performance improvement.

Pilot C: AI-Enabled MRV and Carbon Registry for Distributed Renewables

Scope. Automate MRV for a cluster of municipal rooftop and small utility solar installations and feed verified reductions into a local carbon registry.

Objectives:

- Demonstrate an automated, low-cost MRV pipeline that produces auditable tCO₂e reductions.
- Enable outcome-based finance and pilot issuance of local carbon credits.
- Build trust in AI-verified climate finance.

Core activities:

- Integrate generation telemetry, satellite and imagery cross-checks, and weather forecasts.
- Develop AI models to detect anomalies, estimate baseline counterfactuals, and flag tampering.

- Implement a registry ledger with provenance metadata and third-party audit hooks.
- Pilot issuance of credits to a small buyer pool under strict verification.

Key performance indicators:

- Cost per verification event reduced by at least 50% against manual baseline.
- Time to verify an emission reduction event reduced to days, not months.
- Number of verified tCO₂e and credits issued.
- Third-party audit pass rate of 100%.

5.7 Policy Intervention Summary

The three pilot programmes above constitute the blueprint's policy intervention package, presented here in the Fellowship's standard format. The table is set landscape for legibility.

Table 6.4 Policy intervention summary, in the Capstone Handbook's prescribed format.

Intervention	Description	Implementing Body	Resources Needed	Timeline	Indicators
1. Energy Demand Management (Municipal Building Retrofit)	Retrofit a municipal building with sensors, smart meters and an AI controller to optimise HVAC, lighting and storage, producing verifiable kWh and tCO ₂ e savings for outcome-based contracting.	Lead: Ministry of Works & Housing † Supporting: Ministry of Energy; City & Municipal Governments †	Occupancy sensors; smart thermostats; sub-metering; local edge controller; rooftop solar forecast feed; battery dispatch logic; AI occupancy model	6 months baseline; 12 months active operation; continuous MRV	• 15–25% reduction in annual energy use • 10–20% peak demand reduction • Verified kWh and tCO₂e reported monthly • Occupant comfort at or above baseline
2. Transport Corridor AI Traffic Optimization	Deploy AI signal optimisation across a 2–4 km mixed-traffic urban corridor to cut idling, emissions and travel time, with a public congestion and emissions dashboard.	Lead: Ministry of Works & Housing † Supporting: Transport Agencies; City & Municipal Governments †	Vehicle and pedestrian detection sensors (camera or radar with privacy filters) or existing loop detectors; traffic signal controllers; edge AI optimiser; public dashboard	A/B testing across baseline and optimised weeks; full duration not specified by authors	• ≥20% idling time reduction on corridor • ≥10% average travel time reduction • Reduction in modelled gCO₂e/km • Public transit on-time performance improvement
3. AI-Enabled MRV and Carbon Registry for Distributed Renewables	Automate MRV for a cluster of municipal rooftop and small utility solar installations, feeding verified reductions into a local carbon registry and piloting credit issuance.	Lead: Ministry of Environment (Carbon Registry & MRV Systems) † Supporting: Central Bank; Independent Auditors; Climate MRV Technical Units †	Generation telemetry; satellite and imagery cross-checks; weather forecasts; AI anomaly and counterfactual models; registry ledger with provenance metadata; third-party audit hooks	Not specified by authors	• ≥50% reduction in cost per verification event • Verification time reduced to days, not months • Number of verified tCO₂e and credits issued • 100% third-party audit pass rate

† Implementing bodies marked with a dagger are derived from the authors' own Institutional Mapping (Table 6.7, §6.3); the pilot descriptions do not name a lead ministry directly. Cells marked "not specified by authors" were absent from the submitted text and have not been inferred.

6. Governance and Implementation Model

6.1 The Coordination Layer

Coordination is the most underestimated policy instrument.

Table 6.5 The cost of fragmentation against the coordination mechanisms this blueprint establishes.

Without a coordination layer	This blueprint institutionalizes coordination through
Ministries duplicate efforts	Inter-ministerial steering committees
Cities pilot incompatible systems	Shared data standards
Data standards diverge	Clear mandates and sequencing
AI tools fail to scale	Joint accountability mechanisms

Key insight: Data readiness does not equal implementation readiness.

6.2 National Carbon Registry Architecture

The National Carbon Registry (NCR) is the central institutional and technical mechanism for translating Ubuntuia's AI-enabled urban climate actions into verifiable, finance-grade carbon assets. Its primary purpose is to establish trusted Measurement, Reporting, and Verification (MRV) systems that meet international standards, thereby closing the climate finance gap.

The NCR is structured around a four-layer architecture designed for security, transparency, and scalability:

Table 6.6 National Carbon Registry: four-layer architecture.

Layer	Function	Key Components	Governance / Technology
Layer 1: Data Ingestion	Collects raw, real-time data from urban systems.	Smart meters, IoT sensors, satellite imagery, municipal databases, DRE system logs.	Standardized API endpoints, data governance framework (privacy-preserving), data scaffolding inputs.
Layer 2: AI-Enabled MRV Engine	Transforms raw data into verified emissions reductions.	AI models (forecasting, anomaly detection, fraud detection), emissions baseline models, verification algorithms.	Auditable AI models, Explainable AI Charter compliance, Climate MRV Technical Unit oversight.
Layer 3: Registry Core (Ledger)	Records, tracks, and issues carbon credits.	Secure central database, unique serial numbering system, transaction log.	Central Bank oversight, alignment with CORSIA/VCS/Gold Standard protocols, publicly accessible metadata.
Layer 4: Finance & Reinvestment	Connects verified credits to global finance and directs reinvestment.	Public-facing dashboard, international registry interoperability gateway, Climate Finance Reinvestment Fund.	Ministry of Finance / Central Bank, transparent reinvestment loops.

Key AI functions in Layer 2

- **Continuous emissions measurement.** AI models analyze real-time data to calculate emissions and reductions against a pre-defined baseline.
- **Anomaly and fraud detection.** Machine learning algorithms flag unusual data patterns that may indicate data manipulation or sensor failure.
- **Impact verification.** AI correlates urban interventions with measured emissions reductions to verify the causal link.
- **Public transparency and reinvestment.** A public-facing dashboard (Layer 4) will display the total verified emissions reductions, carbon credits issued and sold, and a transparent breakdown of how the mobilized finance is being reinvested into climate resilience projects (e.g., green jobs, DRE expansion, flood defenses).

6.3 Institutional Mapping (Core Governance)

Clear ownership avoids diffusion of responsibility. Each function has one accountable lead, with clearly defined supporting actors for coordination and execution.

Table 6.7 Institutional mapping: accountable leads and supporting actors by function.

Function	Lead Institution (Accountable)	Supporting Actors (Coordinating & Executing)
Urban Systems & Smart Cities	Ministry of Works & Housing	City and municipal governments; urban planning authorities; transport agencies; private infrastructure operators
Energy & Renewables	Ministry of Energy	National utilities; independent power producers (IPPs); rural electrification agencies; energy regulators
Data, AI & Digital Infrastructure	Ministry of Digital Economy	National statistics office; universities and research institutes; innovation hubs; telecom operators
Carbon Registry & MRV Systems	Ministry of Environment	Central Bank; independent auditors; environmental protection agencies; Climate MRV technical units
Climate Finance & Fiscal Policy	Ministry of Finance	Multilateral climate funds; development banks; budget office; public procurement authorities

Why this table matters:

- One lead per function ensures accountability.
- Supporting actors are coordinated, not competing.
- Prevents policy failure caused by overlapping mandates.
- Enables sequenced implementation rather than fragmented pilots.
- Directly addresses Africa's implementation gap, not just policy design.




Figure 6.3 Government AI and climate policy implementation framework, from cabinet authority to verified carbon credits and capital mobilisation.




Figure 6.4 Decision authority and delivery chain, from strategic policy direction through to implementation.

These institutional flows and authority mappings ensure that coordination precedes implementation. Rather than assuming that clean data or advanced AI systems alone can drive outcomes, this framework recognizes that policy failure in emerging economies is most often a coordination failure. By defining who decides, who executes, and who finances before

deployment, Ubuntia reduces fragmentation risk and increases the likelihood of sustained impact.

6.4 Coordination Mechanism

An Inter-Ministerial AI & Climate Task Force provides:

- Cabinet-level authority
- Mandated data-sharing
- Sequenced implementation
- Conflict resolution

This is the difference between policy and paperwork.

6.5 Implementation Phases

Table 6.8 Four-stage implementation pathway.

Stage	Focus	Core activities
Stage 1	Foundation & Scoping	• Policy problem definition • Institutional alignment • Data audit • Pilot city selection • Stakeholder mapping • Data scaffolding: for data-poor environments, implement low-cost, low-tech data collection methods (e.g., community-led mapping, citizen science, manual data entry via mobile apps) to establish a baseline before deploying high-cost IoT or AI systems, ensuring inclusivity and local ownership of the initial data layer.
Stage 2	Ethics, Risks & Opportunities	• Data governance design • Risk classification • Bias assessment • Public consultation • Explainable AI charter: a charter mandating that public-facing dashboards show not only outcomes (e.g., traffic flow) but also a simple, human-readable explanation for any AI-driven decision that directly impacts a citizen. • Bias red lines
Stage 3	Governance Design	• Coordination frameworks • Legal instruments • Registry architecture • Procurement rules • Data ownership laws • MRV specifications
Stage 4	Integration & Presentation	• System integration • Carbon credit issuance • Policy presentation • Live dashboard showing CO₂ emissions • Jobs created

6.6 Explainable AI Charter

The Explainable AI Charter (XAI Charter) operationalizes the core policy principle: “If I don’t understand it, I will not adopt it.” Its purpose is to ensure that all Artificial Intelligence systems deployed in public services are transparent, accountable, and comprehensible to the citizens they serve.

Scope and applicability

This Charter applies to all High-Risk AI Systems procured, developed, or deployed by government entities. High-Risk AI Systems are defined as those where an erroneous or non-transparent decision could directly impact the fundamental rights or safety of citizens, result in significant financial or environmental loss, or perpetuate systemic bias.

The three levels of explanation

Every High-Risk AI System must provide explanations tailored to three distinct audiences:

Table 6.9 *Explainable AI Charter: the three levels of explanation.*

Level	Audience	Purpose	Explanation Requirement
Level 1: Technical	AI developers, auditors, researchers	System validation and debugging	Model documentation: full technical specifications, training data provenance, model architecture, performance metrics, and bias assessment reports.
Level 2: Operational	System operators, ministry officials	Decision justification and oversight	Counterfactual explanations: for any specific decision, provide the minimum change in input data that would have resulted in a different outcome. Must be auditable by the lead institution.
Level 3: Citizen-Facing	General public, affected individuals	Public trust and right to appeal	Simple, human-readable rationale: a non-technical, concise explanation of why a decision was made, what data was used, and how to appeal. Must be accessible via public dashboards or mobile interfaces.

Mandates for implementation

- **XAI compliance clause.** Mandatory in all government AI contracts.
- **Local mode development.** Preference for systems where explanation mechanisms are trained on and sensitive to local context, dialects, and cultural norms.
- **Mandatory rationale.** For any AI-driven decision that directly affects a citizen, the Level 3 explanation must be provided automatically.
- **Chief AI Transparency Officer (CATO).** Appointed within the Ministry of Digital Economy to audit XAI compliance and serve as the final point of contact for public appeals.
- **Right to explanation.** Citizens have a legally protected Right to Explanation for any adverse decision made by a High-Risk AI System.

Example rationale: “Traffic light timing adjusted due to a 20% increase in pedestrian volume detected by the Kifaru City sensor network, prioritizing safety at the intersection of Unity and Freedom Avenues.”

7. Risk and Mitigation Plan

Table 6.10 Risks and corresponding safeguards.

Risk	Mitigation
Data misuse	Public oversight; privacy-preserving technologies
Institutional conflict	Clear mandates; Inter-Ministerial Task Force
AI opacity	Explainable AI Charter; public-facing decision dashboards
Public distrust	Transparency; communication; public consultation
Energy constraints	AI-Energy Nexus (distributed renewable energy); prioritize low-power AI models
Carbon market crashes	Diversify climate finance sources
Foreign vendor lobbying	Data sovereignty laws; local mode development
Imported bias	Local mode development; bias red lines

8. Monitoring and Evaluation Framework

This framework consolidates the baseline indicators and pilot-level key performance indicators specified elsewhere in this blueprint into a single monitoring structure.

8.1 Baseline Indicators

Table 6.11 Baseline indicators as stated in the blueprint.

Domain	Baseline indicator
Transport	Urban transport emissions growing at 8–12% annually
Energy	Renewable energy share of urban electricity below 30%
Employment	Youth unemployment above 30%, higher for young women
Governance	Over 60% of digital public projects across Africa fail to scale due to weak inter-ministerial coordination
Climate finance	Africa receives less than 5% of global climate finance

8.2 Performance Indicators and Targets

Table 6.12 Consolidated performance indicators and targets, by source pilot.

Domain	Indicator	Target	Source
Building energy	Reduction in annual energy use	15–25%	Pilot A
Peak demand	Reduction in peak demand	10–20%	Pilot A
Verified savings	kWh and tCO ₂ e reported	Monthly	Pilot A
Occupant comfort	Survey score	At or above baseline	Pilot A
Traffic	Idling time reduction on corridor	≥20%	Pilot B
Travel time	Average travel time reduction	≥10%	Pilot B
Vehicle emissions	Reduction in modelled gCO ₂ e/km	Not specified by authors	Pilot B
Public transit	On-time performance improvement	Not specified by authors	Pilot B
MRV cost	Cost per verification event against manual baseline	≥50% reduction	Pilot C
MRV speed	Time to verify an emission reduction event	Days, not months	Pilot C
Carbon credits	Verified tCO ₂ e and credits issued	Not specified by authors	Pilot C
Audit integrity	Third-party audit pass rate	100%	Pilot C

Indicators and targets in this table are drawn verbatim from the pilot programmes in §5.6. Where the submitted text states an indicator without a numeric target, this is recorded as “not specified by authors” rather than inferred.

8.3 Expected Impact

- Faster urban service delivery
- Increased renewable penetration
- Verified emissions reductions
- Access to global carbon markets
- Youth employment in AI and climate services
- Citizens can explain how AI decisions affect them
- Funding for green jobs for unemployed youth
- Strengthened national sovereignty




Figure 6.5 Impact dimensions: environmental, economic, social, governance and technological.

9. Conclusion and Call to Action

Africa must choose: to consume intelligence, or to author it.

This blueprint is an invitation — not to adopt AI blindly, but to govern it wisely, in service of people, planet, and prosperity.

“Technology is neither good, nor bad, nor is it neutral.”
— *Kranzberg’s First Law*

This is not just a capstone. It is a policy artifact.

References

The reference list submitted with this chapter consisted of five short titles without authors, dates, publishers or locators. The volume editors identified the publications those titles refer to and supplied full bibliographic details. Every URL and DOI below was checked and resolved successfully at the time of publication.

- Dener, C., Nii-Aponsah, H., Ghunney, L. E., & Johns, K. D. (2021). *GovTech maturity index: The state of public sector digital transformation*. World Bank. <https://doi.org/10.1596/978-1-4648-1765-6>
- Gold Standard Foundation. (2024). *Principles and requirements* (Version 2.1). Gold Standard Foundation. https://globalgoals.goldstandard.org/standards/101_V2.1_PAR_Principles-Requirements.pdf
- International Civil Aviation Organization. (2023). *Annex 16 to the Convention on International Civil Aviation — Environmental protection, Volume IV: Carbon Offsetting and Reduction Scheme for International Aviation (CORSIA)* (2nd ed.). ICAO. <https://www.icao.int/corsia>
- Meattle, C., Balm, A., Diaz, J. E., Miao, G., Stout, S., Gupta, I., Abraham, S., Sloley, A., & Wignarajah, D. (2024). *Landscape of climate finance in Africa 2024*. Climate Policy Initiative. <https://www.climatepolicyinitiative.org/publication/landscape-of-climate-finance-in-africa-2024/>
- Microsoft AI Economy Institute. (2025). *AI diffusion report: Mapping global AI adoption and innovation* (H1 2025). Microsoft. <https://github.com/microsoft/ai-diffusion-report>
- Misra, A., Wang, J., McCullers, S., White, K., & Lavista Ferres, J. (2025). *Measuring AI diffusion: A population-normalized metric for tracking global AI usage*. arXiv. <https://doi.org/10.48550/arXiv.2511.02781>
- Shehu, M., & Onunwa, G. (2025). *State of AI policy in Africa 2025*. Column Content. <https://columncontent.com/state-of-ai-policy-africa-2025/>
- Verra. (2024). *VCS standard* (Version 4.6). Verra. <https://verra.org/wp-content/uploads/2024/03/VCS-Standard-v4.6.pdf>

Editorial note on sources

Three of the five submitted titles resolved unambiguously to a single publication: the Microsoft AI Diffusion Report, State of AI Policy in Africa 2025, and the World Bank GovTech index. Two required editorial judgement and are recorded here for transparency:

- “World Bank GovTech Index” is formally the GovTech Maturity Index. The 2021 monograph is cited because its DOI resolves reliably; a 2022 update exists but its registered DOI did not resolve at the time of checking.
- “African Development Bank Climate Finance Outlook” does not correspond to any AfDB publication of that title. The Climate Policy Initiative’s Landscape of Climate Finance in Africa 2024 is cited instead, as it is the primary source for the climate finance share the chapter relies on.
- “Global Carbon Market Standards” is not a single publication. It has been expanded into the three standards the chapter names in §6.2 — ICAO CORSIA, the Verified Carbon Standard, and the Gold Standard.

Verification of statistical claims

The submitted text carries no in-text citations. The editors attempted to trace each statistical claim to a primary source. Findings are recorded below; the authors’ wording has not been altered.

Table 6.13 Editorial verification of statistical claims.

Claim as written	Section	Editorial verification
“Africa receives less than 5% of global climate finance”	§3.3	Supported. Climate Policy Initiative (2024) puts Africa’s share of tracked global climate finance below 3%, so the authors’ figure is conservative and defensible.
“Over 60% of digital public projects fail to scale due to weak inter-ministerial coordination”	§3.1	Not verified. No primary source for this figure was located. The underlying qualitative finding — that coordination failure rather than technology is the binding constraint — is well supported in the literature, but the 60% figure could not be traced.
“Urban systems will drive over 70% of future emissions growth in rapidly urbanizing African economies”	§2.3	Partially supported. The ~70% figure in the literature refers to the current global urban share of total emissions, not to the future emissions growth increment in African economies specifically.
“Urban transport emissions growing at 8–12% annually”	§2.2	Illustrative. Presented by the authors as a baseline indicator for a fictional case-study country. Directionally consistent with motorisation trends in fast-growing African cities, but no primary source states this range.
Urban emissions by sector: transport 35%, energy 30%, buildings 20%, waste 15% (Figure 6.2)	§2.2	Un sourced. These proportions appear only in the figure and nowhere in the submitted text. Retained and labelled “illustrative”.

Editorial judgements

No author review cycle was available before publication. Where the submitted material contained defects whose resolution could be established with confidence from the surrounding text, the volume editors have resolved them rather than leaving them in the published record. Each such resolution is listed below. No substantive claim, target, or policy position has been altered.

Table 6.14 Editorial judgements applied to this chapter.

Location	As submitted	Editorial resolution
Working group block	The first author's surname appeared as "Engr Ismail Musa mar".	The letter "U" had been lost to the same font-encoding fault that produced the other text corruption in this file. Restored to "Ismail Musa Umar", consistent with the closing signature block of the submitted document. Honorifics and professional titles are not attached to contributor names anywhere in this volume.
Throughout	Eighteen occurrences of "U" in place of the letter pair "ff" — "diU'usion", "traU'ic", "eU'iciency", "aU'ect", "ScaU'olding", "insuU'icient", "diU'erence", "eU'orts".	Repaired to diffusion, traffic, efficiency, affect, scaffolding, insufficient, difference, efforts.
Throughout	Forty-eight unresolved "fi" and "fl" ligature glyphs.	Expanded to the letter pairs "fi" and "fl". These display correctly on screen but break full-text search, copy-paste and screen readers, so expansion is required for a deposited record.
§1, §2.1, §2.3, §3.2, §9	Five paragraph breaks had been lost, producing run-on sentences such as "...climate finance reinvestment loops At its core lies a guiding principle:".	Breaks restored at the sentence boundaries indicated by capitalisation and sense.
§5.4, §6.1, §6.2, §6.3, §6.6, §7	Five tables had been flattened by PDF conversion into approximately fifty loose, column-major paragraphs, separating headings from their content.	Reconstructed as native tables. Reading order was recovered from the column-major sequence and cross-checked against the screenshots embedded in the submitted file.
§2.3, §5.4, §6.1, §6.5, Table 6.3	Six infographics carried misspellings rendered into the image pixels — including "MISALAGNMENT", "SOVERGENITY", "Institiunallizes", "Standarlized", "REDUTIONS", and "dos nost equal".	Each image was a picture of text already present in the chapter, so each was rebuilt as native content. Nothing was lost, and the material is now searchable and accessible to screen readers.
Figure 6.5	The Impact Dimensions artwork labelled one quadrant "Social" twice, leaving the Governance quadrant mislabelled, and carried a garbled centre mark.	Redrawn using the five dimensions named in the original artwork, populated with the impacts the authors list in §8.3.
Cover artwork	The submitted cover image was cropped at source, clipping the left edge of the title panel so the subtitle began "n Africa-Driven...".	The damaged panel was removed and the chapter now opens with a typographic landing page. No text was lost, since the same title and subtitle appear in full on that page.
§5.7 Policy Intervention Table	Absent. The Capstone Handbook requires interventions in a fixed six-column format.	Built from Pilots A, B and C. Implementing bodies are drawn from the authors' own Institutional Mapping and marked with a dagger; cells with no source in the submitted text read "not specified by authors" and have not been inferred.

Location	As submitted	Editorial resolution
§8 Monitoring and Evaluation	Absent. The chapter listed aspirations under Expected Impact but no indicators, baselines or targets.	Consolidated from indicators the authors had already written inside Pilots A, B and C, plus the baseline indicators in §2.2, §3.1 and §3.3. No indicator or target has been introduced.
References	Five bare titles with no authors, dates, publishers or locators.	Each traced to a real publication and set in APA 7th style, with every URL and DOI checked. Two entries required judgement and are documented in the editorial note above.
Working group block	The roster table carried personal email addresses for both authors.	Withheld. Publishing personal email addresses in a permanent, publicly indexed DOI record is a privacy exposure; institutional affiliation is retained.

Health and Wellbeing

Responsible AI for Health Equity

A Roadmap for Equitable, Trusted Healthcare Where AI Strengthens — Not Replaces — the Human Touch

CASE STUDY

The Republic of Ubutia

WORKING GROUP

Health and Wellbeing Working Group

Samuel Sanaa Brobbey · Ministry of Health, Ghana

Ijeoma Chiemela · VitalNutrient Nutrition Foundation, Nigeria

Genet Shanko Dekebo · Adama Science and Technology University, Ethiopia

Edna C. Too · Chuka University, Kenya

Zubair Usman · Jigawa State Ministry of Health, Dutse, Nigeria

Suggested citation: Brobbey, S. S., Chiemela, I., Dekebo, G. S., Too, E. C., & Usman, Z. (2026). Responsible AI for Health Equity. In N. G. Abdulhamid (Ed.), *Eight Capstones, One Country* (pp. xx–xx). OpenSchool Initiative.
<https://doi.org/10.5281/zenodo.22066741>

1. Executive Summary

Ubuntia stands at a critical juncture in healthcare transformation. With high maternal mortality, fragmented health information systems, and stark urban–rural disparities, the nation faces challenges that artificial intelligence could help address — yet also risks that could deepen existing inequalities if AI is deployed without intentional governance.

This Responsible AI for Health Equity Roadmap presents a comprehensive framework for leveraging AI to strengthen Ubuntia’s healthcare system while safeguarding against algorithmic bias, data exploitation, and technological exclusion. Our vision is clear: deliver equitable, trusted healthcare where AI strengthens — not replaces — the human touch.

Key policy objectives

- Establish interoperable patient data infrastructure that protects privacy while enabling continuity of care
- Develop ethical guidelines for AI-assisted diagnostics that prioritize accuracy, fairness, and clinical oversight
- Enable AI deployment in rural and underserved areas without compromising quality or equity
- Create robust safeguards against algorithmic bias and exclusion in health AI systems

The roadmap proposes four core policy interventions with clear implementation pathways, timelines, and success metrics. Together, these interventions will position Ubuntia as a leader in responsible health AI governance while addressing the immediate healthcare access and quality challenges facing its citizens.

2. Context and Rationale

2.1 Ubuntia's Healthcare Landscape

Ubuntia's healthcare system reflects the nation's broader development patterns: pockets of innovation amid systemic fragmentation. The country's 42 million citizens face a double burden of infectious diseases and rising non-communicable conditions, while maternal mortality remains persistently high. Health information is scattered across disconnected public and private systems, preventing coordinated care and evidence-based policymaking.

The disparities are stark. In Kifarur City and Lumora, private hospitals experiment with AI-assisted triage and digital health records. Meanwhile, a mother in rural Kifarur waits weeks for basic test results, and rural clinics lack even reliable electricity. A single rural teacher manages 70 students with chalkboards while urban peers deploy tablet-based learning — the same pattern repeats in healthcare.

2.2 The AI Opportunity and Risk

AI technologies offer transformative potential for Ubuntia's health system:

- **Clinical decision support.** AI can assist overworked clinicians with diagnostic suggestions, treatment protocols, and patient risk stratification.
- **Health system efficiency.** Predictive analytics can optimize resource allocation, drug inventory management, and appointment scheduling.
- **Remote care extension.** Telephonic health platforms enhanced with AI symptom checkers can extend specialist expertise to rural areas.
- **Disease surveillance.** Pattern recognition algorithms can detect outbreak signals earlier than traditional surveillance systems.

However, without intentional governance, these same technologies could:

- Amplify bias if trained on non-representative datasets that exclude rural populations or women
- Erode privacy as patient data flows through unregulated digital health platforms
- Deepen inequality if AI tools remain confined to wealthy urban hospitals
- Displace healthcare workers without creating alternative opportunities
- Create dependency on foreign technology companies with misaligned incentives

2.3 Why Governance Matters Now

Ubuntia's 2021 Data Protection Act signals commitment to digital rights, but enforcement capacity remains thin. Multiple AI health pilots operate with no unified oversight framework. Without proactive policy intervention, market dynamics will determine AI deployment patterns, likely reinforcing existing inequalities rather than addressing them.

The country faces a choice: become a testing ground for technologies designed elsewhere, or shape AI governance that reflects Ubuntu values of Ubuntu, our shared humanity. This roadmap chooses the latter path.

3. Problem Analysis

Our analysis identifies four interconnected challenges that AI governance must address.

3.1 Data Fragmentation and Interoperability Gaps

Ubuntia's health data exists in silos. Public hospitals use paper records or incompatible electronic systems. Private facilities guard patient data as competitive assets. Laboratory results do not automatically reach ordering physicians. Patients carry physical files between providers, risking loss and duplication of costly tests.

This fragmentation prevents AI systems from accessing the comprehensive patient histories needed for accurate predictions. It also blocks population health analytics that could guide resource allocation and disease surveillance.

3.2 Absence of Ethical Frameworks for Health AI

No national guidelines govern AI use in clinical settings. Diagnostic AI tools enter the market with minimal validation using Ubuntian patient populations. Questions remain unanswered: Who is liable when AI provides incorrect clinical suggestions? How should AI-assisted decisions be documented? What level of clinical oversight is required? When should patients be informed that AI influenced their care?

Without clear ethical guardrails, healthcare providers face legal uncertainty while patients lack transparency about algorithmic decision-making affecting their health.

3.3 Urban–Rural Technology Divide

AI health innovations cluster in Kifaru City and Lumora where reliable electricity, internet connectivity, and technical expertise exist. Rural areas — where 60% of Ubuntians live — face structural barriers:

- **Infrastructure.** Inconsistent power and limited internet access (48% national penetration, far lower in rural areas).
- **Workforce.** Few rural clinicians have digital health literacy or AI training.
- **Economics.** High implementation costs and low patient volumes make rural AI deployment financially unattractive.
- **Language.** Most AI systems operate in English or French, excluding populations who primarily speak Ubuntu or local languages.

Market forces alone will not close this divide. Policy intervention is essential to ensure equitable access.

3.4 Algorithmic Bias and Exclusion Risks

Most health AI systems entering Ubuntia are trained on datasets from high-income countries with different disease patterns, genetic diversity, and healthcare delivery models. This creates multiple bias risks:

- **Diagnostic accuracy.** Algorithms may perform poorly on skin tones underrepresented in training data, affecting dermatology and imaging applications.
- **Treatment recommendations.** Clinical decision support tools may reflect treatment protocols inappropriate for Ubuntia's healthcare infrastructure or drug availability.

- **Resource allocation.** Predictive models could systematically disadvantage marginalized groups if trained on data reflecting historical discrimination.
- **Gender disparities.** Women's health conditions may be underdiagnosed if AI systems are trained predominantly on male patient data.

Without mandatory bias testing and local validation, AI systems risk encoding and amplifying existing health inequities.

4. Policy Objectives

This roadmap pursues four interconnected objectives.

Table 7.1 The four policy objectives of the roadmap.

Objective	Commitment
Objective 1 Data Infrastructure	Establish a national health information exchange that enables secure, patient-controlled data sharing across providers while maintaining strict privacy protections and regulatory compliance with the 2021 Data Protection Act.
Objective 2 Ethical AI Standards	Create comprehensive ethical guidelines for health AI that ensure clinical oversight, mandate bias testing, establish liability frameworks, and require transparent disclosure of AI involvement in clinical decision-making.
Objective 3 Equitable Access	Deploy AI health solutions in rural and underserved areas through infrastructure investment, workforce training, and sustainable financing mechanisms that prevent the technology divide from deepening health disparities.
Objective 4 Bias Prevention	Implement mandatory validation requirements for health AI systems using representative Ubuntian datasets, with ongoing monitoring to detect and mitigate algorithmic bias before it harms patients.

5. Proposed Policy Interventions

The roadmap proposes four policy interventions. Each is presented in the Capstone Handbook's standard format, with description, implementing bodies, resources, success indicators, and timeline.

5.1 National Health Information Exchange (NHIE)

Table 7.2 Intervention 1: National Health Information Exchange (NHIE).

Element	Detail
Description	Establish a federated health information exchange infrastructure that enables secure sharing of patient data across public and private healthcare facilities. The NHIE will use international interoperability standards (FHIR/HL7) while incorporating patient consent mechanisms and robust cybersecurity protections. The system will support AI applications by providing anonymized datasets for research and model training while protecting individual privacy through differential privacy techniques.
Implementing Body	Lead: Ministry of Health, Digital Health Unit Supporting: Ministry of Technology & Digital Innovation; Data Protection Commission; National Health Insurance Authority; Association of Private Hospitals; Medical and Dental Council
Resources Needed	Financial: \$15–20 million initial investment (infrastructure, software licensing, implementation) Human: 25–30 technical staff (system architects, data security specialists, interoperability engineers, training coordinators) Technical: Cloud infrastructure, FHIR-compliant middleware, identity management systems, secure APIs
Success Indicators	• 80% of public hospitals connected to the NHIE within 18 months • 50% of private facilities participating within 24 months • Average patient record retrieval time under 5 seconds • Zero reported data breaches in the first 24 months of operation • At least 3 AI research projects using anonymized NHIE datasets by month 30
Timeline	• Months 1–6: Standards development, regulatory framework, procurement • Months 7–12: Pilot with 5 hospitals (3 public, 2 private) in Lumora • Months 13–18: National rollout to all tertiary and secondary facilities • Months 19–24: Extension to primary care centres and rural clinics • Month 24+: Ongoing expansion, maintenance, and feature enhancement

5.2 Health AI Ethics and Validation Framework

Table 7.3 *Intervention 2: Health AI Ethics and Validation Framework.*

Element	Detail
Description	Establish a mandatory certification process for all AI systems used in clinical decision-making in Ubutia. This framework will require prospective validation using Ubutian patient data, bias testing across demographic subgroups, clinical safety assessments, and regular post-deployment monitoring. The framework includes clear liability allocation, clinical oversight requirements, patient disclosure standards, and appeals mechanisms for contested algorithmic decisions. All certified systems will be listed in a public registry with performance metrics.
Implementing Body	Lead: Medical and Dental Council (regulatory authority expansion) Supporting: Ministry of Health (policy); National Health Research Ethics Committee; Association of Medical Specialists; Patient Rights Advocacy Coalition; Technology Standards Bureau
Resources Needed	Financial: \$3–5 million annually (regulatory capacity building, validation infrastructure, monitoring systems) Human: 15–20 specialized staff (clinical AI evaluators, bioethicists, medical statisticians, patient advocates) Technical: Test datasets representing Ubutian population diversity, validation protocols, monitoring dashboards, public registry platform
Success Indicators	• Ethics framework adopted and published by month 6 • At least 10 health AI systems undergo validation by month 12 • Zero certified systems showing statistically significant demographic bias by month 18 • Public registry accessible with transparent performance data by month 9 • 95% of clinicians report understanding the liability framework by month 24
Timeline	• Months 1–3: Multi-stakeholder consultations and framework drafting • Months 4–6: Framework finalization, regulatory approval, public comment period • Months 7–9: Validation infrastructure setup, evaluator training, registry development • Month 10: Framework enforcement begins, grace period for existing systems • Months 10–24: Rolling certification, monitoring, framework refinement

5.3 Rural AI Health Access Initiative

Table 7.4 Intervention 3: Rural AI Health Access Initiative.

Element	Detail
Description	Launch a comprehensive programme to deploy AI-enabled health solutions in rural and underserved areas, addressing infrastructure, workforce, and financing barriers. The initiative will install solar-powered health IT infrastructure in 100 priority rural clinics; train 500 rural health workers in digital health literacy and AI-assisted care delivery; deploy offline-capable AI tools for symptom assessment, diagnostic support, and patient education in local languages; and establish a sustainable financing model combining government subsidies, health insurance coverage, and results-based donor funding.
Implementing Body	Lead: Ministry of Health, Primary Healthcare Department Supporting: Rural Electrification Agency; Ministry of Education (health training institutions); National Health Insurance Authority; Development Partners Forum; Community Health Workers Association; Telecommunications Regulatory Commission
Resources Needed	Financial: \$8–12 million over 3 years (infrastructure \$5M, training \$2M, software \$1.5M, operations \$3.5M) Human: 50 master trainers, 500 rural health workers, 20 programme coordinators, 10 technical support staff Technical: Solar power systems, ruggedized tablets, offline-capable AI applications, satellite internet backup, multilingual interfaces (English, French, Ubuntu and 3 major local languages)
Success Indicators	• 100 rural clinics equipped with AI infrastructure by month 18 • 500 health workers trained and actively using AI tools by month 20 • 30% reduction in unnecessary patient transfers to urban facilities by month 24 • Diagnostic accuracy in participating rural clinics within 10% of urban benchmarks by month 30 • 80% patient satisfaction with AI-assisted care in rural settings by month 24 • Zero widening of urban–rural health outcome gaps during implementation
Timeline	• Months 1–4: Site selection, infrastructure assessment, technology procurement • Months 5–8: Infrastructure installation (power, connectivity, hardware) in first 20 clinics • Months 9–12: Health worker training programme launch, pilot in initial sites • Months 13–18: Expansion to a further 80 clinics, continuous training • Months 19–24: Full operation, monitoring, impact evaluation • Months 25–36: Scale-up based on learnings, sustainability transition

5.4 Algorithmic Bias Prevention and Monitoring System

Table 7.5 *Intervention 4: Algorithmic Bias Prevention and Monitoring System.*

Element	Detail
Description	Establish a national capability for bias prevention comprising: representative health datasets reflecting Ubuntu's demographic, geographic, and clinical diversity for AI validation; standardized bias testing protocols across gender, age, ethnicity, language, socioeconomic status, and geographic location; real-time monitoring infrastructure to detect bias signals in deployed systems; rapid response mechanisms for bias remediation; and public transparency reporting on AI system performance across population subgroups. The system will collaborate with regional and international bias research networks while maintaining Ubuntuian data sovereignty.
Implementing Body	Lead: National Health Research Institute, AI Governance Division (new unit) Supporting: Ministry of Health; Data Protection Commission; National Statistics Bureau; academic medical centres (University of Lumora, Kifaru Medical College); Civil Society Health Justice Coalition; Gender Ministry
Resources Needed	Financial: \$4–6 million over 3 years (dataset curation \$1.5M, monitoring infrastructure \$2M, research capacity \$1.5M, operations \$1M) Human: 25–30 specialists (data scientists, biostatisticians, health equity researchers, ethics advisors, community engagement coordinators) Technical: Secure data repository, bias detection algorithms, performance monitoring dashboards, public reporting platform, collaborative research infrastructure
Success Indicators	• Representative validation datasets available for 10 major disease categories by month 12 • Standardized bias testing protocols adopted as national standard by month 9 • 100% of newly certified health AI systems pass bias testing by month 18 • Real-time monitoring active for all deployed clinical AI systems by month 24 • Public bias reports published quarterly from month 12 • At least 5 bias issues detected and remediated before patient harm by month 30 • Zero persistent unexplained performance gaps across demographic groups by month 36
Timeline	• Months 1–6: Dataset collection methodology, ethics approval, community consent processes • Months 7–12: Dataset curation from diverse clinical sites and populations • Months 7–9: Bias testing protocol development through a multi-stakeholder process • Months 10–15: Monitoring infrastructure development and deployment • Month 12: First datasets available for validation testing • Months 12–24: Continuous expansion of datasets, refinement of monitoring • Month 12+: Quarterly public reporting begins and continues

6. Governance and Implementation Model

6.1 Institutional Architecture

Successful implementation requires coordinated action across multiple institutions:

- **Health AI Governance Council.** Multi-stakeholder oversight body chaired by the Ministry of Health, including clinicians, ethicists, technologists, patient advocates, and civil society. Meets quarterly to review progress, address emerging issues, and guide policy evolution.
- **Technical Implementation Unit.** Within the Ministry of Health, coordinates the NHIE rollout and rural access initiative. Staffed with project managers, engineers, and health informaticians.
- **Regulatory Capacity.** Expanded Medical and Dental Council with a dedicated AI evaluation unit for certification and ongoing monitoring.
- **Research and Monitoring.** National Health Research Institute hosts bias prevention and monitoring functions, building long-term scientific capacity.




Figure 7.1 Health AI governance architecture, from the multi-stakeholder Governance Council through delivery bodies to stakeholder engagement.

6.2 Stakeholder Engagement

- **Healthcare providers.** Co-design implementation approaches, provide clinical expertise for validation, deliver frontline training.
- **Technology companies.** Comply with certification requirements, share technical documentation, participate in bias testing, invest in local language and context adaptation.

- **Patients and communities.** Participate in consent processes for data sharing, provide feedback on AI-assisted care experiences, hold institutions accountable through advocacy organizations.
- **Academic institutions.** Conduct validation research, train the health AI workforce, develop contextually appropriate algorithms, publish findings to inform policy evolution.
- **Development partners.** Provide catalytic funding, share global best practices, support regional collaboration, respect Ubuntuian policy leadership.

6.3 Financing Strategy

Total estimated investment: \$30–43 million over three years. The proposed financing mix is as follows.




Figure 7.2 Financing strategy: proposed funding mix across four sources.

6.4 Coordination Mechanisms

- Quarterly Health AI Governance Council meetings for strategic oversight
- Monthly technical working group sessions across intervention streams
- Bi-annual public consultations to ensure transparency and accountability
- Annual impact evaluation with independent external review
- Regional collaboration through the Pan-African Economic Bloc digital health working group

7. Risk and Mitigation Plan

Table 7.6 *Risks, likelihood, and mitigation measures.*

Risk	Likelihood	Description and mitigation
Implementation Delays	Medium	Description: Procurement bottlenecks, bureaucratic inefficiencies and infrastructure challenges slow rollout. Mitigation: Dedicated fast-track procurement authority for health AI initiatives; embedded project management unit with clear accountability; regular progress monitoring with escalation protocols.
Technology Company Resistance	Medium	Description: Vendors oppose certification requirements, withdraw from the Ubuntu market, or lobby for weaker standards. Mitigation: Engage companies early in framework design; demonstrate certification benefits (market differentiation, reduced liability); coordinate with regional regulators for harmonized standards; develop domestic AI capacity to reduce foreign technology dependency.
Privacy Breaches	High	Description: Data breaches or unauthorized access undermine public trust in the health information exchange. Mitigation: Defence-in-depth cybersecurity architecture; mandatory security audits; breach response protocols; cyber insurance; patient notification systems; strong legal penalties for violations; independent Data Protection Commission oversight.
Insufficient Funding	Medium	Description: Budget constraints or donor withdrawal reduce available resources below critical thresholds. Mitigation: Phased implementation allowing priority focus; diverse funding sources to reduce dependency; cost-sharing with the private sector; efficiency measures; advocacy for sustained political commitment; clear return-on-investment demonstration to unlock additional investment.
Rural Infrastructure Barriers	High	Description: Persistent electricity and connectivity challenges prevent rural AI deployment. Mitigation: Solar power with battery backup; offline-capable AI applications; periodic connectivity for updates; mobile network partnerships; gradual expansion tied to infrastructure readiness; realistic expectations in remote areas.
Workforce Resistance	Low	Description: Clinicians fear AI displacement or reject new workflows. Mitigation: Frame AI as clinical support, not replacement; involve clinicians in design and validation; demonstrate efficiency gains; provide comprehensive training; protect jobs during transition; celebrate success stories; address liability concerns clearly.
Bias Persistence	Medium	Description: Despite safeguards, algorithmic bias causes patient harm. Mitigation: Continuous monitoring; rapid remediation protocols; transparent reporting; affected patient support; regular bias audits; diverse datasets; inclusive development teams; strong enforcement of certification requirements; public accountability.
Political Will Erosion	Medium	Description: Leadership changes or competing priorities reduce policy commitment. Mitigation: Multi-stakeholder ownership beyond government; institutionalization in permanent structures; demonstrated quick wins; public engagement creating political demand; alignment with broader health system goals; regional commitments through the Pan-African Economic Bloc.

8. Monitoring and Evaluation Framework

8.1 Key Performance Indicators

Table 7.7 Key performance indicators by category.

Category	Indicator	Target
System-level	Healthcare facilities connected to the NHIE	80% of public hospitals by month 18
System-level	Health AI systems certified through the ethics framework	10 by month 12
System-level	Rural clinics with operational AI infrastructure	100 by month 18
System-level	Documented instances of bias detected and remediated	5 or more by month 30
Health outcome	Maternal mortality in areas with AI-assisted prenatal monitoring	Reduction
Health outcome	Diagnostic accuracy in rural facilities using AI decision support	Improvement
Health outcome	Unnecessary patient transfers from rural to urban facilities	Reduction
Health outcome	Urban–rural gap in selected health outcomes	Should not widen; ideally narrow
Equity	AI diagnostic accuracy parity across demographic groups	Parity across gender, ethnicity, age, geography
Equity	Access to AI-assisted care across income quintiles	Equitable access
Equity	Patient satisfaction with AI-assisted care	Disaggregated by demographic group
Equity	Representation in validation datasets	Matches national population distribution
Governance	Reported privacy breaches	Zero
Governance	Clinicians trained in health AI governance	95% by month 24
Governance	Public awareness of AI use in healthcare settings	Measured through surveys
Governance	Stakeholder satisfaction with governance processes	Measured through annual survey

8.2 Data Collection Methods

- **Administrative data.** NHIE usage logs, certification records, training completion rates.
- **Clinical data.** Health outcomes from electronic health records, diagnostic accuracy metrics.
- **Surveys.** Patient experience, clinician feedback, public awareness studies.
- **Qualitative assessments.** Focus groups, case studies, stakeholder interviews.
- **Independent audits.** Technical performance reviews, bias testing results, security assessments.

8.3 Reporting and Accountability

- **Quarterly progress reports.** Submitted to the Health AI Governance Council and published online.
- **Annual impact evaluation.** Independent external evaluation with public presentation.
- **Public dashboard.** Real-time KPI tracking accessible to all stakeholders.
- **Parliamentary briefings.** Annual presentation to the Health Committee with question opportunity.
- **Community feedback mechanisms.** Hotline and online platform for citizen input and complaints.

9. Conclusion and Call to Action

Ubutia stands at a defining moment. The choices made today about health AI governance will shape health outcomes for decades and signal whether technology serves equity or deepens division.

This Responsible AI for Health Equity Roadmap offers a clear path forward — one grounded in African values of Ubuntu, informed by global best practices, and tailored to Ubutian realities. The four interventions are ambitious yet achievable, addressing the full ecosystem from data infrastructure to ethical oversight, from rural access to bias prevention.

Success requires sustained political commitment, adequate financing, multi-stakeholder collaboration, and public trust. It demands that we move beyond technology solutionism to ask harder questions: AI for whom? AI decided by whom? AI that benefits whom?

9.1 What Makes This Different

This roadmap rejects the false choice between innovation and equity. It demonstrates that responsible AI governance can accelerate health system transformation while protecting vulnerable populations. It positions Ubutia not as a passive technology consumer but as an active policy shaper — influencing regional standards, contributing to global debates, and building domestic capacity for the AI era.

Most importantly, this roadmap keeps the human touch at the centre. AI should strengthen — not replace — the compassion, clinical judgment, and cultural understanding that make healthcare healing. Technology that distances providers from patients or erodes trust will ultimately fail, no matter how sophisticated its algorithms.

9.2 Immediate Next Steps

1. Cabinet endorsement of the roadmap and allocation of initial funding in the next budget cycle
2. Establishment of the Health AI Governance Council with diverse stakeholder representation
3. Launch of multi-stakeholder consultations to refine implementation details
4. Designation of lead implementing agencies with clear mandates and accountability
5. Development of detailed operational plans for each intervention with monthly milestones
6. Initiation of funding mobilization from government, health insurance, and development partners

9.3 The Path Forward

The mother in rural Kifaru who waits weeks for test results deserves better. The farmer in the Central Highlands threatened by erratic rainfall and food insecurity deserves access to AI-enhanced preventive care. The young medical graduate in Lumora deserves to practice in a system where technology augments rather than threatens professional dignity.

This roadmap is their roadmap. It translates the promise of health AI into concrete action. It transforms the abstract goal of “responsible innovation” into measurable interventions with clear accountability.

Ubutia can lead Africa in demonstrating that AI governance is not a barrier to innovation but its foundation; that equity and excellence are not trade-offs but mutually reinforcing goals; that technology can honour our shared humanity rather than erode it.

*The time to act is now. The roadmap is clear. The future of health equity in Ubutia
— and across Africa — depends on the choices we make today.*

References

The submitted chapter did not include a reference list. The volume editors have identified authoritative sources for the external frameworks and standards the chapter relies on. Every URL and DOI below was checked and resolved successfully at the time of publication.

Gichoya, J. W., Banerjee, I., Bhimireddy, A. R., Burns, J. L., Celi, L. A., Chen, L.-C., Correa, R., Dullerud, N., Ghassemi, M., Huang, S.-C., Kuo, P.-C., Lungren, M. P., Palmer, L. J., Price, B. J., Purkayastha, S., Pyrros, A. T., Oakden-Rayner, L., Okechukwu, C., Seyyed-Kalantari, L., ... Zhang, H. (2022). AI recognition of patient race in medical imaging: A modelling study. *The Lancet Digital Health*, 4(6), e406–e414. [https://doi.org/10.1016/S2589-7500\(22\)00063-2](https://doi.org/10.1016/S2589-7500(22)00063-2)

Health Level Seven International. (2019). *HL7 FHIR Release 4*. HL7 International. <https://hl7.org/fhir/R4/>

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>

World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance*. World Health Organization. <https://www.who.int/publications/i/item/9789240029200>

World Health Organization. (2024). *Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models*. World Health Organization. <https://www.who.int/publications/i/item/9789240084759>

Editorial note on sources

No references were submitted with this chapter. The sources above were located by the volume editors to support the external standards and evidence the chapter invokes: the FHIR interoperability standard named in Intervention 1, the WHO guidance underpinning the ethics framework in Intervention 2, and the peer-reviewed literature on algorithmic bias in health systems that motivates Intervention 4. They are offered as supporting apparatus; the authors' analysis and wording are unchanged.

Verification of statistical and legal claims

Table 7.8 Editorial verification of claims.

Claim as written	Section	Editorial verification
“Ubuntia’s 2021 Data Protection Act”	§2.3	Internal to the case study. Ubuntia is the Fellowship’s fictional country, so this instrument is part of the scenario rather than a real statute.
“42 million citizens”	§2.1	Internal to the case study, and consistent with the population figure used across other chapters in this volume.
“60% of Ubuntians live in rural areas”	§3.3	Internal to the case study, and consistent with the figure used in Thematic Area 6.
“48% national internet penetration”	§3.3	Internal to the case study. No external source is cited; the figure is presented as a scenario parameter.
Algorithmic bias risks in diagnostic imaging and resource allocation	§3.4	Externally supported. Obermeyer et al. (2019) and Gichoya et al. (2022) document both effects in deployed health systems. Citations supplied by the editors.
All financial figures (\$15–20M, \$3–5M, \$8–12M, \$4–6M, \$30–43M total)	§5, §6.3	Authors’ own estimates for the case study. Presented as proposed budgets rather than sourced costings.

Editorial judgements

No author review cycle was available before publication. Where the submitted material contained defects whose resolution could be established with confidence from the surrounding text, the volume editors have resolved them rather than leaving them in the published record. Each such resolution is listed below. No substantive claim, target, budget, or policy position has been altered.

Table 7.9 Editorial judgements applied to this chapter.

Location	As submitted	Editorial resolution
§5.1 to §5.4, Interventions 1–4	Each intervention table had been split across two or three separate tables. Rows were severed mid-sentence: Intervention 2’s Implementing Body ended “... Association of” and resumed in a different table with “Medical Specialists, Patient Rights Advocacy Coalition”. Intervention 4 had lost its Description label entirely.	Each intervention reassembled into a single table in the Handbook’s standard format, with the severed sentences rejoined and the missing row labels restored.
§7 Risk and Mitigation Plan	The risk table was fragmented across three separate tables, with the “Bias Persistence” row split so that its mitigation appeared in a different table beneath an empty row label.	Reassembled into a single eight-row table with likelihood and mitigation columns intact.
§7, all rows	Every mitigation cell began “Mitigation: Mitigation:”.	The duplicated label was removed and mitigation text placed in its own column.
§8.1 and §8.2	Indicators appeared as five loose bulleted groups — system-level, health outcome, equity, governance — without a common structure.	Consolidated into a single indicator table retaining the authors’ own categories and targets. No indicator or target has been introduced.
§4 Policy Objectives	The four objectives ran as unlabelled prose paragraphs.	Presented in tabular form under the authors’ own objective names. Wording unchanged.
§6.1 and §6.3	The institutional architecture and financing mix were described in prose only; the chapter contained no figures.	Two diagrams were drawn from that text — the governance architecture and the 40/30/25/5 financing split — so the structures are legible at a glance. Both restate content already present in the chapter.
References	Absent. The chapter cited external standards without a reference list.	Five sources located and verified by the editors for the standards the chapter invokes — FHIR, WHO guidance on AI in health, and the peer-reviewed literature on algorithmic bias underpinning Intervention 4. Documented in the editorial note above.
Working group block	The roster table carried personal email addresses for all five authors.	Withheld. Publishing personal email addresses in a permanent, publicly indexed DOI record is a privacy exposure; institutional affiliation is retained.

Chapter information

Table 7.10 Chapter information as recorded by the authors.

Field	Detail
Project	AI Literacy Fellowship for African Policymakers
Working Group	Health and Wellbeing
Capstone Output	Responsible AI for Health Equity Roadmap
Country Context	The Republic of Ubuntia
Period	October 2025 – February 2026
Organizing Institution	OpenSchool Initiative and Cosmopolitan University

This blueprint is part of the integrated National AI Policy Blueprint for Ubuntia: Governing AI for Equity, Innovation, and Sovereignty.

Justice, Rights and Governance

A Justice, Rights and Governance AI Framework

Positioning Ubuntu as an Active Rule-Setter Rather Than a Passive Consumer of Externally Designed Technology

CASE STUDY

The Republic of Ubuntu

WORKING GROUP

Justice, Rights and Governance Working Group

Abubakar Usman Abubakar · Universiti Malaya, Malaysia

Muhammad Nasambo · Nigeria Revenue Service, Nigeria

Oluwaseyi Opeyemi Jimoh · Osun State Ministry of Innovation, Science, Technology and Digital Economy, Nigeria

Sadiya Ibrahim Galadima · National Judicial Institute, Nigeria

Suggested citation: Abubakar, A. U., Nasambo, M., Jimoh, O. O., & Galadima, S. I. (2026). A Justice, Rights and Governance AI Framework. In N. G. Abdulhamid (Ed.), *Eight Capstones, One Country* (pp. xx–xx). OpenSchool Initiative. <https://doi.org/10.5281/zenodo.22066741>

1. Executive Summary

Artificial intelligence (AI) is rapidly reshaping governance, service delivery, and socio-economic development across the world. While AI offers significant opportunities to address longstanding development challenges, its unregulated adoption risks deepening inequality, enabling rights violations, and reinforcing external technological dependency — particularly in developing and postcolonial contexts. This capstone proposes a Justice, Rights, and Governance AI Framework for the Republic of Ubuntia, a Sub-Saharan African state characterized by demographic diversity, uneven development, and growing digital adoption.

The framework responds to Ubuntia’s increasing reliance on externally designed, donor-driven AI systems that often lack local relevance, accountability, and safeguards. Without a coherent national governance approach, AI deployment risks reproducing algorithmic bias, exclusion of rural and marginalized populations, data misuse, and elite capture. These challenges are compounded by institutional capacity constraints, procurement vulnerabilities, and persistent digital divides.

To address these risks, the proposed framework advances a rights-based, inclusive, and accountable model of AI governance grounded in democratic values and local realities. It establishes robust institutional structures — including a National AI Governance Council and an independent AI Ethics and Rights Commission — supported by comprehensive legal instruments on data protection, algorithmic accountability, and public-sector AI procurement. The framework prioritizes local capacity-building, gender-responsive and rural-inclusive AI deployment, transparency, and meaningful public participation.

A phased implementation strategy and a strong monitoring and evaluation framework ensure adaptability, accountability, and long-term sustainability. By positioning Ubuntia as an active rule-setter rather than a passive technology consumer, the framework demonstrates how African states can govern AI in ways that promote justice, protect rights, and advance inclusive development. Ultimately, this capstone argues that equitable AI governance is not merely a technical challenge but a strategic choice about democratic futures in the digital age.

2. Context and Rationale

Over the past decade, artificial intelligence has emerged as a transformative force across nearly all sectors of human activity, including education, healthcare, economic development, governance, and everyday social interactions. AI-driven systems are reshaping how societies generate knowledge, deliver services, and make decisions. While these technologies offer significant benefits — such as improved efficiency, innovation, and expanded access — they also generate serious risks, including algorithmic bias, social exclusion, surveillance, and the deepening of existing inequalities. These dynamics underscore the urgent need for governance frameworks that ensure AI advances justice, protects rights, and strengthens accountability.

Across Africa, AI adoption is accelerating, yet concerns are growing that unregulated or poorly governed deployment may reinforce structural inequalities rather than alleviate them. Sub-Saharan Africa faces persistent political, economic, and developmental challenges that demand innovative and scalable policy responses. AI holds considerable potential to support solutions in sectors such as agriculture, healthcare, education, and public administration — but only if guided by inclusive, ethical, and rights-based governance frameworks.

2.1 The Ubuntu Context

Within this context, the Republic of Ubuntu is a Sub-Saharan African country marked by a youthful population, rich cultural diversity, and a relatively young democratic system. The country is linguistically, ethnically, and religiously plural, reflecting broader regional patterns. At the same time, Ubuntu faces interconnected development challenges. Over sixty percent of the population resides in rural areas, driving sustained rural–urban migration to cities such as Lumora and Kifaru. Agriculture employs more than half of the workforce but remains low in productivity. Educational access is uneven, youth unemployment exceeds thirty percent, and women face disproportionately high barriers to economic and social mobility.

Despite these constraints, Ubuntu has achieved moderate digital connectivity. Approximately 48 percent of the population has internet access, supported by relatively reliable 4G coverage and widespread social media use. AI technologies have already entered the national landscape through pilot projects in agriculture, healthcare, and financial services. However, most initiatives are donor-funded and rely on imported AI systems that inadequately reflect local languages, cultural contexts, and indigenous data. This raises concerns about technological dependency, data extraction, and the reproduction of external norms and priorities.

These conditions highlight the need for a national AI strategy explicitly grounded in justice, rights, and inclusive governance. Without deliberate policy intervention, Ubuntu risks becoming a testing ground for externally designed technologies that do not serve local public interests. Although progress has been made in digital service delivery at the central government level, local governments remain under-resourced, producing uneven access and widening governance gaps. A carefully designed AI governance framework — rooted in transparency, accountability, and social inclusion — can help harness AI for public good while preventing the deepening of inequality.

3. Problem Analysis

AI adoption in the Republic of Ubuntia is increasing, yet it remains fragmented, weakly regulated, and largely externally driven. The absence of a coherent national AI governance framework exposes the country to significant risks, including rights violations, elite capture of technological benefits, and the amplification of existing inequalities. As AI systems increasingly shape public service delivery and social life, the lack of safeguards threatens both social justice and democratic accountability.

Institutional weaknesses compound these risks. Regulatory oversight of emerging technologies is underdeveloped, particularly at subnational levels where administrative capacity is limited. While anti-corruption initiatives exist, procurement vulnerabilities and accountability gaps persist, raising concerns about transparency in AI acquisition and deployment.

Ubuntia's reliance on donor-funded and foreign-designed AI systems further entrenches dependency. These systems often lack local-language datasets and cultural relevance, limiting inclusivity and effectiveness. More critically, they position Ubuntia as a consumer rather than a rule-setter in the global AI ecosystem.

Socioeconomic inequality and digital exclusion exacerbate these challenges. Internet access and AI-enabled services remain concentrated in urban areas, while rural populations, women, and unemployed youth face systematic exclusion. In this context, unregulated AI adoption is more likely to reproduce inequality than promote inclusive development.

Without effective governance, AI systems risk reinforcing algorithmic discrimination, enabling surveillance and data misuse, and undermining fundamental rights such as equality, transparency, and access to redress. If left unaddressed, these dynamics threaten public trust, democratic legitimacy, and long-term development. A comprehensive, rights-based AI governance framework is therefore essential.

4. Policy Objectives

The overarching objective of the Justice, Rights, and Governance AI Framework is to ensure that AI development and deployment in Ubutia advance inclusive, equitable, and sustainable development while safeguarding fundamental rights and democratic governance. Specifically, the framework seeks to:

1. Promote rights-based and ethical AI use
2. Strengthen accountable and transparent AI governance
3. Reduce inequality and digital exclusion
4. Build local capacity and reduce external dependency
5. Enhance public trust and democratic participation

5. Proposed Interventions

5.1 Overview and Strategic Objectives

This intervention establishes a comprehensive, rights-based AI governance framework that positions Ubuntu as an active rule-setter rather than a passive recipient of externally designed technologies. Grounded in justice, accountability, transparency, and social inclusion, the framework responds to Ubuntu's socioeconomic realities and institutional constraints while advancing long-term development goals.

The intervention pursues four strategic objectives: to develop rights-protective legal and policy instruments; to establish robust institutional and regulatory foundations; to build local AI capacity and reduce technological dependency; and to promote inclusive and equitable AI deployment.



Figure 8.1 The four strategic objectives of the intervention.

5.2 Institutional Architecture

- **AI Ethics and Rights Commission (AERC).** An independent body with investigative and enforcement powers to oversee compliance, assess high-risk AI systems, investigate discrimination complaints, and issue binding recommendations.
- **National AI Council (NAIC).** A multistakeholder body coordinating AI strategy, ensuring cross-sectoral alignment, and publishing annual public reports. Capacity-building programmes and regional AI focal points support subnational governance.

5.3 Legal and Regulatory Framework

Table 8.1 Proposed legal and regulatory instruments.

Instrument	Provisions
Data Protection and Privacy Act	Establishes individual data rights, consent requirements, enforcement mechanisms, and protections for vulnerable and illiterate populations.
AI Procurement and Public Sector Standards	Mandates transparent procurement, bias assessments, local-language support, technology transfer, and third-party audits.
Algorithmic Accountability and Transparency Act	Requires impact assessments, public registries of high-risk systems, transparency obligations, and legal avenues for contesting automated decisions.

5.4 Local AI Capacity Development

- National AI Research Institute focused on local-language models and development priorities
- Ubuntu Data Commons Initiative ensuring ethical, community-owned datasets
- AI Skills Development Programmes targeting policymakers, youth, and women
- Innovation support mechanisms including sandboxes, seed funding, and incubation

5.5 Inclusive and Equitable Deployment

- Rural digital infrastructure expansion
- Mandatory algorithmic fairness audits
- Public awareness and digital literacy campaigns
- Gender-responsive AI policy and youth engagement initiatives

5.6 Summary of the Intervention

This intervention presents a comprehensive and context-sensitive framework for AI governance in the Republic of Ubuntu that balances ambition with pragmatism. It recognizes both AI's transformative potential to address development challenges and the serious risks it poses to justice, rights, and equity if left ungoverned. By establishing robust institutions, enacting rights-protective legislation, investing in local capacity, and centering inclusion, the framework positions AI as a tool for public good rather than a driver of inequality.

Crucially, the framework rejects the assumption that African states must passively adopt externally designed technologies and governance models. Instead, it affirms Ubuntu's agency as a rule-setter capable of developing AI governance rooted in its own values, priorities, and social contexts. Through deliberate, rights-based policy choices, Ubuntu can demonstrate that AI governance in African contexts can be inclusive, accountable, and socially beneficial — offering lessons for the wider region and global community.

While success will require sustained political commitment, adequate resourcing, and genuine public participation, the alternative — fragmented, externally driven AI adoption — poses far greater risks to democratic governance and social cohesion. This intervention therefore represents not only a technical policy framework, but a strategic choice about Ubuntu's technological future.

6. Governance and Implementation Model

The governance and implementation model is designed to ensure that AI systems in Ubuntuia are deployed in ways that are inclusive, rights-respecting, and responsive to local contexts.




Figure 8.2 Governance and implementation architecture, from the apex oversight bodies through delivery mechanisms to the legal foundation.

6.1 National AI Governance Council

Composition. The NAIGC will include representatives from government ministries, civil society organizations, youth groups, technology startups, academia, and community leaders, ensuring broad societal representation.

Responsibilities. The Council will oversee implementation of the national AI strategy, coordinate across sectors, monitor compliance with ethical and rights-based standards, and assess impacts on marginalized populations.

6.2 Public-Private Partnerships

Collaboration. PPPs will leverage private-sector expertise while embedding public-interest safeguards. Partnerships will prioritize culturally relevant, context-sensitive AI solutions, particularly in sectors such as agriculture and healthcare.

Joint initiatives. Pilot projects will be co-designed with local stakeholders to ensure relevance, accountability, and scalability.

6.3 Decentralized Implementation

Local empowerment. Local governments will be equipped with authority and resources to implement AI solutions tailored to community needs, particularly in rural and underserved areas.

Capacity building. Targeted training programmes will strengthen AI literacy and governance capacity among local officials and community actors.

6.4 Community Engagement Mechanisms

Public consultations. Regular consultations will gather citizen feedback on AI initiatives, ensuring transparency and inclusion.

Feedback and redress. Accessible mechanisms will allow citizens to report bias, exclusion, or harm, enabling continuous improvement and trust-building.

7. Risk and Mitigation Plan

Implementing a justice- and rights-based AI governance framework in Ubuntu will encounter political, institutional, and structural constraints. This framework anticipates these challenges and embeds mitigation strategies within its design.

Table 8.2 Risks and mitigation strategies.

Risk	Description	Mitigation
Corruption and procurement irregularities	Corrupt procurement may result in inappropriate or harmful AI systems, and procurement vulnerabilities persist alongside existing anti-corruption initiatives.	Public disclosure of AI procurement contracts; independent third-party audits; civil society participation in oversight; whistleblower protection. The AERC retains independent authority to investigate procurement decisions raising rights or accountability concerns.
Surveillance and misuse of AI	AI may be deployed for unjustified surveillance or data misuse.	Strict legal limits, proportionality requirements, and oversight by the AERC.
Institutional and resource capacity constraints	Weak administrative and fiscal capacity may undermine implementation, particularly at subnational level.	Phased approach prioritizing foundational legal and institutional investments; partnerships with universities, regional organizations and international agencies while maintaining national ownership; training-of-trainers models to create multiplier effects; mentorship programmes and sustained training investment.
Algorithmic bias and discrimination	AI systems may reproduce or intensify existing social biases, resulting in discrimination against marginalized groups, particularly women, rural populations, and linguistic minorities.	Mandate regular algorithmic audits to detect and address bias in datasets, models, and outcomes, with AERC oversight; require inclusion of diverse stakeholder groups — especially women, rural communities, and minority language speakers — in the design, testing, and evaluation of AI systems.
Political economy and elite capture	Elite interests may dominate AI governance and capture its benefits.	Multistakeholder institutional design; public consultation requirements; independent regulatory authority; regular public reporting, parliamentary oversight, and access to information mechanisms.
Technological dependency	Continued reliance on externally designed, donor-driven systems entrenches dependency and external norms.	Technology transfer requirements in procurement contracts; preferential use of open-source systems; regional cooperation with other African states to enhance collective bargaining power with global technology providers.
Digital divide and infrastructure constraints	Inadequate digital infrastructure, particularly in rural and underserved areas, may limit equitable access to AI-enabled services.	Explicit rural service delivery targets; investment in broadband, community technology hubs and affordable access initiatives; mobile, offline and community-based AI service delivery models that ensure access where traditional infrastructure is weak or unavailable and reduce reliance on individual digital access; monitoring of whether AI systems reduce or exacerbate inequality.

This table consolidates the authors' risk analysis, which appeared in the submitted text in two places — “Addressing Implementation Challenges” and “Risks and Mitigation Plan”. The two treatments overlapped substantially; they have been merged without adding new risks or mitigations.

By proactively identifying these risks and embedding targeted mitigation strategies, Ubutia can ensure that its AI governance framework advances justice, accountability, and equitable development rather than reinforcing existing inequalities.

8. Monitoring and Evaluation Framework

The Monitoring and Evaluation framework is designed to assess the effectiveness, equity, and accountability of AI governance in Ubutia, ensure compliance with justice and rights principles, and support evidence-based policy adaptation.

8.1 Institutional Arrangements

Table 8.3 Monitoring and evaluation: institutional arrangements.

Element	Detail
Lead Institution	National AI Governance Council, or a designated coordinating ministry
Supporting Institutions	AI Ethics and Rights Commission; Data Protection Authority; Ministry of Justice; Anti-Corruption Commission; sectoral ministries; local governments; civil society organizations; academic institutions
Reporting Timeline	Annual monitoring reports; mid-term review at Year 3; final evaluation at Year 5

8.2 Key Performance Areas and Indicators

Table 8.4 Key performance areas and indicators.

Performance area	Indicators
Governance and regulatory effectiveness	- Adoption, enforcement, and compliance with national AI governance regulations - Transparency and integrity of AI procurement and contracting processes
Rights protection and accountability	- Accessibility and effectiveness of grievance and redress mechanisms - Frequency and outcomes of independent audits addressing bias, discrimination, and misuse - Compliance with data protection, consent, and privacy requirements
Equity and inclusion	- Distribution of AI-enabled public services across rural and urban areas - Use of local languages and culturally relevant datasets - Gender- and youth-disaggregated outcomes in AI-supported services
Capacity development and local innovation	- Public investment in AI skills development and institutional capacity - Growth of local AI startups, research partnerships, and innovation hubs - Reduction in reliance on externally controlled or donor-driven AI systems
Public trust and participation	- Public awareness of AI use in government services - Frequency and inclusiveness of public consultations - Trends in citizen trust in digital governance and public service delivery

The submitted text specifies these indicators without baselines, numeric targets, or reporting frequencies. None have been inferred by the editors; the reporting timeline in §8.1 is the only cadence the authors state.

8.3 Data Collection Methods

- Government administrative and procurement records
- Independent audits, evaluations, and algorithmic impact assessments
- National surveys, stakeholder interviews, and community consultations
- Reports from civil society, academia, and oversight institutions

8.4 Learning, Transparency, and Adaptation

Monitoring and evaluation findings will be publicly disclosed to promote transparency and accountability. Lessons learned will inform adaptive policy adjustments, ensuring that AI governance remains responsive to emerging risks, technological changes, and societal needs.

9. Conclusion and Call to Action

The Republic of Ubuntia stands at a pivotal moment in shaping its technological future. AI presents powerful opportunities to address longstanding development challenges, but without deliberate governance, it also poses serious risks to justice, rights, and democratic accountability. This framework offers a clear pathway for ensuring that AI serves the public interest rather than entrenched power or external dependency.

9.1 Call to Action

6. **Government commitment.** Prioritize the adoption and implementation of the Justice, Rights, and Governance AI Framework in alignment with national development goals and human rights obligations.
7. **Public awareness and participation.** Engage citizens through inclusive public education campaigns, empowering communities to understand AI systems and assert their rights.
8. **Investment in capacity building.** Allocate sustained resources to develop local technical, legal, and institutional expertise, fostering domestic innovation and long-term autonomy.
9. **International and regional cooperation.** Pursue strategic partnerships that support knowledge exchange while preserving national ownership and policy autonomy.

By acting decisively, Ubuntia can demonstrate that AI governance in African contexts need not replicate global inequalities, but can instead advance inclusive, rights-based, and socially grounded innovation. The choices made today will shape whether AI deepens existing divides or becomes a cornerstone of a just and equitable digital future for all citizens.

References

The submitted chapter did not include a reference list. The volume editors have identified authoritative sources for the international instruments and scholarship the chapter engages. Every URL and DOI below was checked and resolved successfully at the time of publication.

- African Union. (2024). *Continental artificial intelligence strategy: Harnessing AI for Africa's development and prosperity*. African Union.
<https://au.int/en/documents/20240809/continental-artificial-intelligence-strategy>
- Birhane, A. (2020). Algorithmic colonization of Africa. *SCRIPTed*, 17(2), 389–409.
<https://doi.org/10.2966/scrip.170220.389>
- Council of Europe. (2024). *Framework convention on artificial intelligence and human rights, democracy and the rule of law* (CETS No. 225). Council of Europe.
<https://www.coe.int/en/web/conventions/full-list?module=treaty-detail&treatynum=225>
- Mohamed, S., Png, M.-T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33, 659–684.
<https://doi.org/10.1007/s13347-020-00405-8>
- Organisation for Economic Co-operation and Development. (2019). *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449). OECD.
<https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. UNESCO.
<https://unesdoc.unesco.org/ark:/48223/pf0000381137>
- United Nations. (2011). *Guiding principles on business and human rights*. United Nations.
<https://www.ohchr.org/en/publications/reference-publications/guiding-principles-business-and-human-rights>

Editorial note on sources

No references were submitted with this chapter. The sources above were located by the volume editors to support the framework's intellectual commitments: the UNESCO Recommendation and OECD AI Principles, which are the leading international instruments on rights-based AI governance; the Council of Europe Framework Convention, the first binding treaty on AI and human rights; the African Union Continental AI Strategy, the regional instrument most directly relevant to the chapter's argument; the UN Guiding Principles on Business and Human Rights, which underpin the procurement and technology-transfer provisions; and the scholarship of Birhane and of Mohamed, Png and Isaac on algorithmic colonialism, which corresponds closely to the chapter's critique of technological dependency and its rule-setter argument. They are offered as supporting apparatus; the authors' analysis and wording are unchanged.

Verification of claims

Table 8.5 Editorial verification of claims.

Claim as written	Section	Editorial verification
“Over sixty percent of the population resides in rural areas”	§2.1	Internal to the case study, and consistent with the figure used in Thematic Areas 6 and 7.
“Approximately 48 percent of the population has internet access”	§2.1	Internal to the case study, and consistent with the figure used in Thematic Areas 4 and 7.
“Youth unemployment exceeds thirty percent”	§2.1	Internal to the case study, and consistent with Thematic Area 6.
“Agriculture employs more than half of the workforce”	§2.1	Internal to the case study. No external source is cited.
Algorithmic bias, surveillance and dependency risks in African AI adoption	§3, §7	Externally supported. Birhane (2020) and Mohamed, Png and Isaac (2020) set out the algorithmic colonialism critique the chapter advances. Citations supplied by the editors.
Rights-based AI governance as an international norm	§1, §5	Externally supported by the UNESCO Recommendation (2021), the OECD AI Principles (2019), and the Council of Europe Framework Convention (2024). Citations supplied by the editors.

Editorial judgements

No author review cycle was available before publication. Where the submitted material contained defects whose resolution could be established with confidence from the surrounding text, the volume editors have resolved them rather than leaving them in the published record. Each such resolution is listed below. No substantive claim, argument, or policy position has been altered.

Table 8.6 Editorial judgements applied to this chapter.

Location	As submitted	Editorial resolution
Section numbering	Section numbers were typed by hand and had lost their anchor. Headings one to seven carried no numbers at all, after which “8.1”, “8.2”, “8.3” and “9.1” appeared — so the first numbered subsection in the chapter was 8.1.	Numbering rebuilt on Word heading styles and applied consistently from §1 to §9. Section order is unchanged.
§5 and §7	Risk was treated twice. “Addressing Implementation Challenges” in the interventions section covered corruption and procurement, resource and capacity constraints, elite capture, technological dependency and the digital divide; the later “Risks and Mitigation Plan” covered substantially the same ground in different words, so a reader met several risks twice.	The two treatments were merged into a single seven-row table in §7. Every risk and every mitigation is the authors’ own; where the two passages described the same risk, the fuller description was retained and the mitigations combined.
§7, formatting	The first three risks ran “Risk: ... Mitigation: ...” together in single paragraphs; the last two used a separate “Mitigation Strategies:” heading with bullets.	All risks now share one structure: risk, description, mitigation.
§5.6 and §9	The chapter contained two sections titled “Conclusion” — one closing the interventions section, one closing the document.	The first was retitled “Summary of the Intervention”, which is what it is. The second remains the chapter conclusion. No text was changed.
§5.1 and §6	The institutional architecture and the four strategic objectives were described in prose only; the chapter contained no figures.	Two diagrams were drawn from that text — the governance architecture and the four strategic objectives — so the structures are legible at a glance. Both restate content already present in the chapter.
§5.3 and §8.1	The legal instruments and the monitoring arrangements were set out as bulleted lists mixing headings with detail.	Presented in tabular form under the authors’ own labels. Wording unchanged.
References	Absent. The chapter engaged international norms and a body of scholarship without a reference list.	Seven sources located and verified by the editors, including the scholarship on algorithmic colonialism that corresponds most directly to the chapter’s rule-setter argument. Documented in the editorial note above.
Working group block	The roster table carried personal email addresses for all four authors.	Withheld. Publishing personal email addresses in a permanent, publicly indexed DOI record is a privacy exposure; institutional affiliation is retained.