

Pax Technologica

DESIGNING AFRICA'S INTELLIGENCE FUTURE:

Essays on AI, Governance,
and Public Value

Selected Essays from Blavatnik School
of Government Graduates, University of Oxford

PEACE

JUSTICE

Pax Technologica

DESIGNING AFRICA'S INTELLIGENCE FUTURE:

Essays on AI, Governance, and Public Value

Selected Essays from Blavatnik School
of Government Graduates, University of Oxford

INTRODUCTION

THOMAS ERMACORA

ESSAYS

MARLON-RALPH NYAKABAU

SIKA SALAMI

NILOFAR SAIYED

ARFAN RAMADHAN

SITONG LYU

OMOTAYO DADA

CONCLUSION

AARON MANIAM

06	INTRODUCTION THOMAS ERMACORA
10	HUMAN-CENTRICITY IN DIGITAL TRANSFORMATION AND AI MARLON-RALPH NYAKABAU
16	BEYOND THE ‘USER’: Recentring Public Value over User-Centred Design in Africa’s Digital Transformation SIAKA SALAMI
26	DATA AS DESTINY: Unlocking Efficiency and Trust in Public Service NILOFAR SAIYED
32	THE RAILS OF INCLUSION: Building Africa’s Digital Foundations for AI ARFAN RAMADHAN
38	SMART CITIES, SMART CONTRACTS: Negotiating AI Partnerships on African Terms SITONG LYU
46	HOW AI CAN ACCELERATE INTRA-AFRICAN TRADE OMOTAYO DADA
54	CONCLUSION AARON MANIAM

THOMAS ERMACORA

For anyone concerned with Africa's future, one question increasingly sits at the centre of the conversation about artificial intelligence. Will AI act as an accelerant—intensifying existing inequalities, constraints, and dependencies—or could it generate a form of escape velocity from long-standing developmental limits? Framed this way, the issue is not whether AI will matter for the continent, but how it will shape trajectories that are already uneven and contested.

This question carries unusual weight because artificial intelligence is not confined to a single sector or policy domain. It cuts across governance, markets, public services, and cultural life, embedding itself in the everyday functioning of societies. Africa is encountering this transformation under conditions that are often fragile, with limited room for error and little margin for delay. As a result, outcomes will not be determined by technical capability alone. They will depend on choices—about design, governance, and direction—that will shape the underlying conditions within which African societies develop.

From this starting point, it is not difficult to tell a worrying story about how artificial intelligence may unfold across African societies. For many observers, the signs already point in a troubling direction. The forces shaping the AI landscape are large, fast-moving, and often far removed from the continent itself. Decisions about investment, standards, intellectual property, and infrastructure are increasingly made elsewhere, according to priorities that do not necessarily align with local needs or constraints.

This dynamic is reinforced by the broader geopolitical environment in which AI is developing. Artificial intelligence has become a central arena of strategic competition, with resources, capabilities, and influence concentrated primarily in two dominant technological ecosystems. As this competition intensifies, the space for late entrants to shape outcomes on their own terms appears to narrow. For countries with limited bargaining power, choices are frequently framed as matters of alignment rather than design: access in exchange for dependence, speed in exchange for control. Seen from this angle, artificial intelligence looks less like a tool for transformation than a force that may deepen existing asymmetries, accelerating challenges that have long constrained self-determined development.

Yet the worrying story is not the only one that can be told. Alongside it sits the possibility of a different trajectory—one that does not begin from alignment with existing technological ecosystems, but from a more basic question: what kinds of technology are appropriate to specific local conditions? This approach might be described as local source: not a rejection of global technology, but a way of grounding its development and deployment in local needs, capacities, and priorities.

Local source starts from problem identification rather than solution transfer. It asks what challenges are most pressing, what constraints are real rather than assumed, and what forms of technology might meaningfully address them. In practice, this can take many forms. Some efforts focus on foundational infrastructure; others operate within specific systems—such as finance, health, food, or climate resilience—where technology can help address persistent constraints. What matters is not the sector, but the alignment between talent development, institutional frameworks, and the creation of businesses capable of responding to these needs over time.

Under the right conditions, some of these initiatives may develop sufficient momentum to become companies that are genuinely rooted in African contexts—able to generate sustainable value, retain agency, and build a measure of economic sovereignty. Even where outcomes are uneven or limited in scale, their significance would extend beyond immediate impact. They would demonstrate that alternative paths are not only imaginable, but workable, and that technological development need not be confined to models defined elsewhere.

The possibility of a different trajectory, however, should not be mistaken for inevitability. Local source is not something that emerges spontaneously from market forces or technological diffusion alone. It depends on deliberate choices and sustained effort. Without that, the gravitational pull of existing ecosystems, incentives, and power structures is likely to prevail. The conditions required for more self-directed forms of technological development do not assemble themselves; they have to be built, protected, and maintained over time.

In this sense, the idea of an African dream is less about aspiration than about cultivation. Cultivating such a dream means creating the political, institutional, and social space in which alternative futures can be taken seriously and acted upon. This work does not rest with any single group. It requires leaders willing to articulate direction and hold long-term horizons in view; policymakers able to shape rules, incentives, and public institutions in ways that enable rather

than constrain possibility; entrepreneurs and engineers prepared to build under real constraints rather than idealized ones; and an engaged civic body—citizens, communities, cultural actors, and the wider public—whose expectations and participation ultimately shape what is legitimate, durable, and worth sustaining. Artificial intelligence is not the only substance of this dream, but it is a central part of it, precisely because of its reach across so many aspects of social and economic life. Without collective commitment across these domains, the more hopeful story remains just that—a possibility left unrealized.

The significance of timing now comes into sharper focus. Even those who disagree about how artificial intelligence should be governed tend to agree on one point: opportunities of this scale do not remain open indefinitely. The choices made in the coming years will shape not only how AI is adopted, but who benefits from it and on what terms. Once certain paths are taken—through infrastructure, regulation, or institutional design—they become difficult to reverse.

For Africa, the consequences of missing this moment would be substantial. In the absence of strong local capacity and governance, digital systems are more easily captured, misused, or exploited. Economic opportunities that might otherwise be developed domestically are absorbed elsewhere, deepening dependence rather than reducing it. Pressures on employment intensify as productivity gains accrue unevenly. At the same time, the expansion of AI-enabled surveillance raises serious questions about civil liberties and accountability, particularly in contexts where institutional safeguards are already under strain. When African societies lack the ability to shape AI systems for their own needs, they also risk seeing their cultures, values, and lived experiences underrepresented—if not misrepresented—in technologies trained predominantly on data from elsewhere.

At the same time, the potential upside of acting effectively is considerable. Africa's position outside the dominant technological ecosystems creates space to explore alternative approaches—local source models that are grounded in specific contexts, constraints, and values. Such approaches would not only matter for Africa itself. They could offer valuable reference points for many other societies that likewise sit beyond the centres of technological power, and that face similar questions about agency, representation, and inclusion.

It is against this backdrop that this publication has taken shape. We invited a group of contributors from the Blavatnik School of Government who participated in Pax Technologica's Fellowship Programme to share their perspectives on how artificial intelligence is beginning to intersect with questions of governance, public value, and societal change in African contexts, particularly in light of recent advances in generative AI. The aim was to bring together a range of viewpoints shaped by work and study across areas such as infrastructure, public services, markets, and institutional design.

From the outset, the intention was not to produce a comprehensive account or a single, unified framework. The AI landscape is too fluid, and the questions it raises too varied, for that to be a realistic—or even desirable—goal. Instead, the essays that follow approach the field from distinct vantage points, illuminating different pressures, constraints, and possibilities. In doing so, they offer both critical reflections and concrete proposals, without pretending that the field's tensions can be fully resolved.

This publication is therefore offered as an invitation: to slow the conversation down at a moment when speed is often rewarded, and to engage more carefully with both the risks and the opportunities that AI presents. It is intended to support leaders, policymakers, practitioners, and citizens who are being asked to make consequential decisions under conditions of uncertainty, and to contribute to a broader discussion about how African societies might shape their engagement with artificial intelligence in deliberate and purposeful ways.

Artificial intelligence is already beginning to shape the conditions under which societies make decisions about power, value, and responsibility. Those conditions are neither neutral nor settled; they are being actively formed. The pages that follow sit within that moment. They engage questions that are still open, choices that are still being negotiated, and futures that are not yet fixed. What matters now is not only the seriousness with which these questions are held, but the clarity and conviction with which responses are developed.

HUMAN-CENTRICITY IN DIGITAL TRANSFORMATION AND AI

MARLON-RALPH NYAKABAU

AFRICA'S ADVANTAGE

Africa holds a unique strategic advantage in the digital era. While the global digital divide has left much of the developing world behind, this gap creates a valuable second-mover opportunity. By observing the successes and shortcomings of early adopters, Africa can harness lessons from first-mover nations where rapid technological advancement has brought both positive and negative socio-economic consequences. These insights can offer a more strategic roadmap for Africa's digital transformation and AI adoption.

Africa is currently advancing a continent-wide digital transformation agenda under the African Union's Digital Transformation Strategy for Africa (2020–2030). Although the continent has historically lagged in the first three industrial revolutions, the onset of the Fourth Industrial Revolution presents a unique opportunity to leapfrog. Within this framework, the African Union identifies digital transformation and artificial intelligence as critical drivers to enhance regional integration, foster inclusive economic growth, create sustainable employment opportunities, and accelerate poverty reduction.

Technologically advanced nations in the Americas, Europe, and Asia clearly demonstrate the benefits of digital transformation. In these regions, innovation has created new industries, raised productivity, lifted millions out of poverty, and improved quality of life. At the same time, these countries also illustrate the double-edged nature of technology. Rapid advancements have brought unintended consequences: widening income inequality, the rise of trillion-dollar oligopolies, deindustrialization, and the proliferation of misinformation and disinformation that erodes trust in public institutions.

For Africa, the challenge is how to capture the opportunities of digitalization and artificial intelligence while mitigating their risks. The continent's comparatively weaker institutions make it particularly vulnerable. Issues such as vaccine hesitancy, tribalism, xenophobia, and radicalization could be amplified by digital platforms, even as technology offers tools to address them. It is therefore imperative that African policymakers design and enforce strategies that maximize the developmental benefits of digital transformation and AI while proactively guarding against its potential harms.

The way forward lies in adopting a holistic perspective on the impact of technology on human lives. By prioritizing human-centricity, African countries can design and implement digital transformation and AI strategies that not only drive innovation and economic growth but also safeguard social cohesion, equity, and well-being.

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

WHAT IS HUMAN-CENTRICITY AND WHY DOES IT MATTER?

There is no single, universally accepted definition of human-centricity in technology. The prevailing view among researchers is that human-centricity emphasizes technology that prioritizes people, aligning with fundamental values such as ethics, rights, justice, dignity, diversity, inclusion, and privacy.

Governments also interpret human-centricity in varied ways. The European Union's AI guidelines emphasize trustworthy, transparent, and accountable systems that advance societal and environmental well-being. Australia's ethical AI framework focuses on respect for human rights, diversity, and individual autonomy. Meanwhile, China's AI governance principles stress the promotion of human-machine harmony and the enhancement of collective well-being.

In practice, the most common application of human-centricity has been user-centric design. This approach treats users as active contributors rather than unpredictable resources to be automated away. Developers engage users throughout the design process, ensuring that technologies meet their needs, goals, and challenges. Such engagement leads to greater usability, satisfaction, and retention, while also fostering inclusivity. However, in user-centricity, there remains a gap between how developers and users perceive human-centricity. Developers prioritize an 'avoidance of harm' approach—while users prioritize the broader social impacts of technology. Genuine human-centric technology requires a stronger emphasis on people's lived experiences and needs.

This stronger emphasis calls for a reimagining of the human-technology relationship. In his book Human-Centered AI, Professor Ben Shneiderman describes human-centricity as the pursuit of technologies that elevate human experience, satisfaction, and performance. Rather than replacing people, such systems aim to amplify, augment, and empower—making technology reliable, safe, and trustworthy, while encouraging creativity, clarifying responsibility, and strengthening social participation. Ultimately, the goal of human-centricity is not to diminish humanity, but to harness technology in ways that expand human potential and improve collective well-being.

HUMAN-CENTRIC DEVELOPMENT

A human-centric perspective also requires a clear and effective understanding of universal human needs. Nobel Laureate Amartya Sen's capability approach offers such a reimagining of development. Sen argued that poverty and well-being should be assessed not only through material resources, but through people's capabilities and freedoms—their real opportunities to live the lives they value. This contrasts with narrower definitions of development that equate progress solely with GDP growth, industrialization, or technological advancement. From this lens, development is equally about the removal of 'unfreedoms' that restrict choice, dignity, and human experience. Extending this logic, a human-centric approach to technology focuses on expanding human capabilities and freedoms, rather than limiting success to the fulfilment of basic needs or material gains.

AFROCENTRICITY

Afrocentricity can be understood as a contextualized expression of human-centricity, with Ubuntu at its core. Ubuntu is an African philosophy of personhood that views the community as the foundation of identity and life. Ubuntu highlights interconnectedness, community, and shared humanity, offering important insights for shaping technology that is inclusive, ethical, and socially beneficial. Within the context of Africa’s digital transformation, Afrocentricity expressed through Ubuntu can serve as a guiding principle for designing technologies that prioritize collective well-being and social harmony over narrow individual gain.

ECONO-CENTRICITY VS HUMAN-CENTRICITY

Digitalization and AI adoption have the potential to transform the way economies, governments, and societies function. Yet in many countries, the primary motivation for digital transformation is framed in economic terms—efficiency, productivity, profit, growth, and competitiveness. This reflects the dominance of techno-economic agendas and econo-centric capitalism, which judge innovation narrowly by economic returns rather than human benefit, often undervaluing social, cultural, and ecological concerns. This paradigm views unlimited economic expansion, assisted by modern technology, as the most effective path to addressing social and developmental challenges.

Such an approach misses the opportunity to harness technology for social and community needs. Human-centricity shifts the focus to what people truly need in their lives and how technology either enables or hinders the fulfilment of those needs. This extends beyond economic functions into areas such as culture, sport, and entertainment, while also supporting projects that foster care, reduce loneliness, and enhance human connection with nature. By contrast, econo-centric approaches risk narrowing technology’s role and may reinforce silos within governments, when social priorities remain under-invested.

A human-centric approach also pays greater attention to the adverse and unintended consequences of technology for society, culture, and the environment. Importantly, the inclusion of women, the elderly, persons with disabilities, and minority groups in the digital economy must not translate into greater exposure to digital harms. Yet if monitoring and evaluation of digitalization and AI investments focuses exclusively on economic metrics—profit, productivity, exports—risks such as polarization, cybercrime, online abuse, misinformation, and disinformation will remain overlooked. These negative externalities must be systematically tracked and addressed.

Unlike econo-centricity, human-centricity is anchored in values and culture. Emerging technologies, including AI, can be harnessed not only to drive innovation but also to preserve and promote the cultural richness and diversity of societies. Efforts to support indigenous content and local technology companies should extend beyond the economic rationale of reducing imports or boosting exports. They should also aim to strengthen indigenous value systems and safeguard traditional cultural heritage, ensuring that digital transformation contributes to both human well-being and cultural continuity.

**CASE STUDY FOR HUMAN-CENTRICITY
SINGAPORE**

Singapore’s digital transformation journey is embodied in the Smart Nation initiative, launched in 2014 with the vision of embedding digital technologies across government, the economy, and society to enable citizens to live meaningful, connected, and fulfilling lives.

A decade later, Singapore’s digital economy has become a powerful growth engine, contributing 17.7 percent of GDP. Yet the initiative extends beyond productivity and competitiveness. From the outset, it sought to enhance both personal and community well-being through technology.

The updated Smart Nation 2.0 outlines several humancentric use cases:

Leveraging AI to monitor and protect trees in parks and neighbourhoods.

Tracking physical activity and rewarding healthy behaviours.

Supporting seniors to remain active and respond to emergencies.

Using AI to assist children’s learning in schools.

Gamifying civic participation through online platforms.

Deploying AI-enabled multimedia in libraries to expand creativity and imagination.

Using digital tools to strengthen engagement with Singapore’s diverse arts, culture, and heritage.

Singaporeans value the ‘kampung spirit’—a deep-rooted sense of community and cooperation. To uphold this ethos in a digital era, the government developed platforms enabling social workers to mobilize goods and services for vulnerable households.

At the same time, Smart Nation recognizes the risks of digitalization. It promotes cyber wellness in schools, encourages screen-free activities for children, and provides support for victims of cyberbullying, ensuring that technological progress does not compromise safety, trust, or social cohesion.

TOWARDS A HUMAN-CENTRIC PARADIGM FOR AFRICA’S DIGITAL TRANSFORMATION AND AI ADOPTION

A paradigm shift in Africa’s approach to digital transformation and AI adoption is essential. Moving away from the prevailing econo-centric focus towards human-centricity will ensure that technology is not only a driver of economic growth but also a tool for enhancing human well-being and community resilience. While user-centricity—as commonly practised—marks a useful step in this transition, it should not be seen as the final goal. Human-centricity goes further, reimagining the relationship between humans and technology with the aim of amplifying and enriching lived experience. As the vision of Pax Technologica suggests, the ultimate aim is to move beyond material gain towards thoughtfully using technology to enhance individual lives and communities.

1. HUMAN-CENTRIC DESIGN AND IDEATION

African policymakers must expand the imagination of technology use cases to reflect the full spectrum of human needs. This means leveraging AI and digital tools not only in business and productivity but also in entertainment, sport, electoral processes, judicial systems, and security services. Such an approach embraces the pluripotential of emerging technologies and ensures they serve broader societal goals.

2. ADDRESSING CROSS-CUTTING DEVELOPMENT CHALLENGES

Technology should also be harnessed to address Africa’s deep-rooted socio-political challenges. As Amartya Sen’s capability approach suggests, development is about removing ‘unfreedoms.’ In this spirit, AI can promote integration and unity by enabling multilingual communication to combat tribalism and xenophobia. Digital tools can increase government transparency to fight corruption, and expand civic engagement to foster peaceful democratic transitions.

3. HUMAN-CENTRIC IMPLEMENTATION

In practice, this involves embedding local languages, arts, and cultural design into technology systems. Doing so preserves African heritage in digital form, strengthens intellectual property rights, and ensures that Africa’s digital transformation is not defined solely by imported Western technologies and values.


4. MONITORING AND EVALUATION

Human-centricity must also guide the assessment of digital policies. Governments should establish watchdog agencies tasked with monitoring the unintended consequences of digitalization, such as mental health impacts, online addiction, hate speech, cyber fraud, and e-waste. Anticipating and addressing these externalities will be crucial for sustainable digital ecosystems.

5. AFROCENTRICITY AS CONTEXTUALIZED HUMAN-CENTRICITY

Finally, Africa can apply human-centricity through Afrocentric frameworks. The philosophy of Ubuntu provides a powerful foundation for participatory design. Embedding Ubuntu in digital transformation means developing tools that reflect indigenous knowledge, communal values, and local traditions. It also suggests promoting open-source platforms that empower African developers to adapt technologies to local needs and cultural contexts.

By shifting towards a human- and Afrocentric model of digital transformation, Africa can ensure that AI and emerging technologies deliver inclusive growth, strengthen social cohesion, preserve cultural identity, and empower communities—rather than replicating econo-centric models that prioritize profit over people.



CONCLUSION

Digital transformation entails leveraging innovation and the adoption of digital products, services, and processes to enhance the functioning of economies, governments, and societies. Yet a review of current digital transformation and AI policies reveals the dominance of a techno-economic agenda rooted in econo-centric capitalism. This paradigm measures technological innovation primarily by its economic returns rather than its contribution to human well-being. It is the same model that fuelled digital transformation in advanced Western economies, resulting in the rise of powerful tech oligopolies and serious negative externalities, disrupting the integrity of politics, health, family life, science, and the law. Should Africa pursue this same path, the consequences could be more severe.

Yet, Africa and the wider developing world now have a unique opportunity not only to catch up, but to lead, by embracing a human-centric model of digital transformation. Such a model, championed by governments, must stand apart from approaches driven primarily by stock market interests. Instead, it should be firmly grounded in meeting human needs and enhancing lived experiences through technology. A human-centric approach reframes digital transformation from a ‘whole-of-government’ or ‘whole-of-economy’ agenda into a ‘whole-of-people’ vision. In Africa, embedding the philosophy of Ubuntu—with its emphasis on community, interconnectedness, and shared humanity—further contextualizes and strengthens this human-centric perspective, ensuring digital transformation delivers inclusive and sustainable benefits for all.

BEYOND THE ‘USER’:

Recentring Public Value over User-Centred Design in Africa’s Digital Transformation

SIAKA SALAMI

INTRODUCTION: WHEN EMPATHY MASKS ONTOLOGY

User-centred design (UCD) promises a humane antidote to bureaucratic abstraction: start from user needs, embrace iterative prototyping, and measure success by satisfaction and uptake. In consumer technology, this ethos has generated elegant devices and compelling interfaces. However, when exported wholesale into the public sector, especially across African states with complex histories of colonial administration, uneven infrastructures, and plural normative orders, UCD risks romanticizing problems it is ill-equipped to solve. Its focus on individual preferences, local usability, and rapid iteration can misrecognize the state’s distinctive obligations: to uphold rights, distribute burdens fairly, steward critical infrastructures, ensure security, and build legitimacy under law. Empathy with ‘the user’ is not the same as justice for the citizen.

This essay advances a philosophically grounded critique of UCD assumptions, compares stylized applications of UCD in the global North and selected African contexts, and evaluates how digital transformation efforts are structured with respect to cybersecurity, AI ethics and governance, and digital skills. It concludes with proposals for ‘public value-centred design’ (PVCD): a normative-institutional alternative that reorients design away from individualized satisfaction towards collective capabilities, procedural fairness, and resilient infrastructures. The provocation is simple: the state is not a start-up, the citizen is not a customer, and delight is not a proxy for dignity.

DESIGNING AFRICA’S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

1. WHAT UCD ASSUMES AND WHY IT OFTEN FAILS GOVERNMENTS

UCD’s origins in human-computer interaction and industrial design foreground an individual decision-maker with stable preferences interacting with a bounded technology experience. This carries five background assumptions:

First, the user is substitutable: if you understand ‘the user’, you can design for the population. However, the public sector serves multiple, sometimes incommensurable publics, including non-users who still bear costs and risks. The citizen denied a social transfer because of a biometric mismatch is a stakeholder, even if she never ‘uses’ the interface.

Second, the value function is experiential: success is measured by convenience, time saved, or satisfaction scores. Government, however, is tasked with solving collective action problems that require equity, due process, transparency, and long-term stewardship, goods not reducible to momentary delight.

Third, participation is intrinsically empowering. UCD methods, such as interviews and co-creation workshops, are considered inclusive. Yet participation can be tyrannical where power asymmetries are unaddressed; selection effects and tokenistic rituals can instrumentalize communities to legitimize pre-made decisions.

Fourth, institutions flex around prototypes. UCD presumes teams can iterate in the wild; bureaucratic systems rarely permit such plasticity because they encode legal mandates, budget cycles, audit trails, and political accountability. Iteration collides with administrative law.

Fifth, infrastructures are dependable and invisible. UCD frequently treats identity, payments, connectivity, language, security, and data governance as exogenous. In fact, these are the very objects of design in public digitization, and they are deeply political.

These assumptions do not doom UCD; they delimit its proper scope. Where governments adopt UCD as a master philosophy, they risk confusing interface empathy for institutional ethics. A passport renewal app that is ‘delightful’ but collapses under a denial-of-service attack, or that quietly excludes citizens without smartphones, is not a public-sector success. Conversely, a plainer system that prioritizes auditability, redundancy, and lawful process may produce lower satisfaction scores while maximizing legitimate public value.

Paul Hughes reflects that: ‘As a designer, I tend to push back against the rigid and fixed definition of user-centred design. My argument is not that user-centred design has not become rigid and misused, but rather that it has become so through the use of non-designers. A designer would never work with user-centred design in isolation from a wider ecology, a wider set of questions, and an ongoing dynamic to check assumptions. A similar, rigid, and fixed definition of a design process takes place with Design Thinking, which has become a five-step process used by non-designers. If designers do indeed work with this process of design thinking, they do so within a much wider, richer context and ecology that surrounds it.’

2. PHILOSOPHICAL LENSES FOR RE-SITUATING DESIGN

A richer ontology of public design in Africa emerges by staging a conversation among four traditions: African communitarian ethics (Ubuntu), the capability approach, post-colonial design theory, and state-capacity realism. A similar sentiment is shared in Marlon’s essay Human-Centricity in Digital Transformation and AI, which explains that ‘Afrocentricity’ expressed through Ubuntu can serve as a guiding principle for designing technologies that prioritize collective well-being and social harmony over narrow individual gain.

Ubuntu rejects atomized individualism by constructing personhood as relational: ‘a person is a person through other persons,’ and privileges restoration, solidarity, and dignity. Designing under Ubuntu reframes ‘needs’ as obligations within a community; the relevant harms include eroding social trust and exacerbating fragmentation. This challenges UCD’s focus on individual convenience by foregrounding collective goods and the moral status of intermediaries who translate state services into lived benefits.

The capability approach, associated with Sen and Nussbaum, centres the fundamental freedoms people have to be and do things they value. Its relevance is twofold. First, services should be judged by capability expansion rather than raw usage. Second, capabilities depend on conversion factors, such as social norms, literacy, and networks, so that identical interfaces can have unequal effects across different contexts. A digital tax system that is ‘usable’ yet requires stable electricity and numeracy may reduce freedom for informal workers by shifting compliance costs onto them. UCD that stops at usability measures the finger, not the moon.

Post-colonial design theory cautions against ‘universal’ design logics that perpetuate coloniality by erasing local epistemologies and life worlds. ‘Design for the pluriverse’ urges designers to treat alternative ontologies as inputs, not obstacles. In African bureaucracies that inherit colonial archives and categories, classification work—i.e., who counts as a citizen and what constitutes a household—must be viewed as design choices with significant political consequences.

State-capacity realism, finally, insists that design must match administrative capability. The ‘Problem-Driven Iterative Adaptation’ (PDIA) tradition emphasizes working from locally nominated problems, experimenting within authorizing space, and building capability through iteration. Problem-Driven Iterative Adaptation complements but constrains UCD: ethnographic empathy without institutional authorization produces prototypes that cannot scale; iterative pilots without political support produce fragile islands of excellence in seas of indifference.

Together, these frames generate three tests for public design: Does it advance collective dignity (Ubuntu)? Does it expand substantive capabilities

(Sen/Nussbaum)? Does it align with state capacity and build it over time (Problem-Driven Iterative Adaptation)? UCD contributes methods for inquiry within these tests, but cannot substitute for them.

3. HOW THE GLOBAL NORTH STRUCTURES ‘USER-CENTRED’ STATES

The global North contains diverse approaches, but several archetypes recur. Estonia’s ‘government as a platform’ vision relies on the X-Road interoperability layer, a national digital ID, and legal norms that let agencies reuse data with transparency logs. The UK Government Digital Service (GDS) has institutionalized service standards, multidisciplinary teams, and a design system that prioritizes simplicity and consistency across services. Legislation, such as the Equality Act, and strong data protection enforcement shape expectations for inclusion and privacy-by-design. In the United States, digital-era governance has oscillated between federal rescue efforts (e.g., the U.S. Digital Service) and agency-led modernization, with a growing emphasis on design standards and platform infrastructure for login, payments, and notifications.

In all three, UCD methods are embedded within capable legal-bureaucratic scaffolds: reliable civil registration, high broadband penetration, robust identity infrastructures, and enforceable privacy regimes. Prototyping tends to sit atop established infrastructures (identity, payment, and address systems), rather than within their construction. When high-profile failures occur, such as the early Healthcare.gov crash, they are often due to capacity mismatches rather than deep ontological questions about the citizen-state relationship.

Crucially, ‘user needs’ in these contexts are framed within rights-bearing legal subjects. Consent management, redress mechanisms, and impact assessments are enforced through courts and regulators; privacy-by-design expectations and cybersecurity baselines shepherd UCD. The state can therefore license experimentation precisely because harms are bounded by institutions. The North’s most successful public design stories, beneath the surface, are tales of governance as much as empathy.

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

4. AFRICAN TRAJECTORIES: INFRASTRUCTURING THE STATE WHILE DESIGNING SERVICES

African governments are pursuing ambitious digitization: national ID systems, service portals and one-stop centres, e-procurement, digital tax and customs, and data-exchange layers. Rwanda’s Irembo has become a canonical example of platform-style service delivery with iterative improvement; Kenya’s eCitizen and Huduma centres represent hybrid digital-physical access; Nigeria’s Treasury Single Account, bank verification number (BVN), and national identification number (NIN) reforms illustrate the infrastructural nature of ‘service design.’ These efforts are shaped by coalitions of domestic reformers, multilateral lenders, and vendors, each carrying their own imaginaries of the ‘modern’ state.

Two structural facts distinguish many African trajectories from Northern cases. First, digitization often proceeds by constructing the very infrastructures of identity, address systems, secure data centres, and broadband backbones that Northern UCD takes as given. That makes ‘the interface’ the least of the problems. The pluralities of identity (names, languages, kinship structures) and the legacies of exclusion (statelessness at borders, nomadic communities) render registration neither neutral nor purely technical. Second, administrations contend with capability traps: pay rigidities, high turnover, procurement rules tuned to concrete rather than code, and vendor lock-in.

UCD methods are visible, including user interviews in informal settlements, rapid prototyping of mobile workflows, and co-creation workshops with civil society. However, they are often layered atop unresolved questions of authority and law. Consider national digital ID projects: they promise inclusive access and fraud reduction, yet when biometric deduplication fails or grievance redress is weak, apparently ‘minor’ usability issues become rights-affecting exclusions from pensions or schooling. Over-indexing on ‘user journeys’ while

under-specifying legal bases for data processing and third-party oversight yields ethically fragile architectures.

Moreover, African publics rarely align with the atomised ‘user.’ Peter Ekeh’s classic essay describes African political life as split between a ‘civic public,’ associated with the state and its formal obligations, and a ‘primordial public,’ rooted in kinship and moral community. Citizens navigate both, often with justified ambivalence towards state power given colonial and postcolonial histories. UCD rituals of empathy can, paradoxically, depoliticize this ambivalence by translating it into ‘pain points’ rather than grappling with legitimacy deficits and historical trauma.

5. THE LIMITS OF EMPATHY: PARTICIPATION, LEGITIMACY, AND THE ‘TWO PUBLICS’

Empathy interviews and co-design workshops risk becoming extraction rituals if they do not shift agenda-setting power. Arnstein’s ladder reminds us that there is a spectrum from manipulation to citizen control. Cooke and Kothari warn of the ‘tyranny’ of participation when elite framing, closed-ended choices, and implicit coercion limit authentic voice. In African digital transformation, donor-funded timelines and procurement constraints can lock in solutions before community deliberation, inviting performative consultation rather than shared governance.

A deeper problem is legitimacy. UCD’s success in commercial settings stems from an implicit social contract: users can exit, and competition disciplines poor design. Citizens cannot exit the state. When trust is low, ‘delightful’ interactions do little to repair the deficit without parallel investments in procedural justice of clear rules, explainability, grievance redress, and visible accountability. Recasting research from empathizing with users to diagnosing legitimacy deficits reframes the work: the target of change becomes institutions, not just interfaces.

Ekeh’s two publics complicate consent. A grandmother may consent to biometric capture to access a pension in the civic public while trusting her primordial public (her family network) to manage risks she cannot fully understand. Contextual integrity theory suggests that privacy norms are domain-specific; violations occur when information flows breach role expectations. Capturing fingerprints in a clinic may be acceptable; reusing them for policing without an explicit warrant may not. UCD’s notion of ‘informed consent’ can be procedurally compliant yet substantively hollow if it ignores the fractured moral topography of consent.

6. CYBERSECURITY: WHEN USABLE DESIGN MEETS ADVERSARIAL REALITY

Public digital systems are vulnerable to fraud, sabotage, and geopolitically motivated intrusion. Classic UCD priorities—reduce steps, eliminate friction, reuse passwords—tend to degrade security; meanwhile, security practices that ignore human factors produce shadow workarounds. In African contexts where SIM-swap fraud, social engineering, and insider threats are prevalent, ‘user-friendly’ authentication can be catastrophic.

A security-by-design posture requires adversarial personas in the same fidelity as user personas: the fraudster who tailgates in queues to harvest identity numbers, the corrupt insider with privileged database access, and the political actor who pressures an agency to exfiltrate opposition records. Threat modelling must be paired with procurement constraints and skills realities: multi-factor authentication using hardware tokens may be unrealistic at scale; SMS-based OTPs are vulnerable; USSD flows may need time-outs tuned for low-end devices and shared phones. Socio-economic realities must contextualize security controls; otherwise, we design brittle defences or exclusionary barriers.

UCD contributes when it illuminates the human-security interface, such as how to stage second-factor prompts without excluding low-literacy users, how to design lockdown policies that mitigate brute-force attacks without denying services to legitimate users, and how to craft recovery processes that resist social engineering. Nevertheless, the metric cannot be satisfaction alone; it must be risk-adjusted assurance. A passport renewal that takes two minutes longer but reduces account takeover by half is a win for public value. If UCD resists such trade-offs, it becomes an anti-security ideology.

7. AI ETHICS AND GOVERNANCE: SHIFTING FROM ‘SMART SERVICES’ TO ACCOUNTABLE SYSTEMS

AI systems already animate public services: biometric deduplication for IDs, risk scoring for audits, triage in social protection, chatbots for information requests, and computer vision in border control. The ethics literature has consolidated principles of beneficence, non-maleficence, fairness, accountability, transparency, and privacy. However, principles without institutions produce ethics theatre. African states face two distinctive governance challenges.

First, data scarcity and bias: training data often reflect Northern contexts or narrow domestic populations, producing models that underperform for linguistic, cultural, or phenotypic diversity. A face-matching model tuned on other populations may amplify exclusion in ID enrolment; a language model that misses idioms in Hausa or isiZulu may misclassify citizen intents in chatbots. Second, procurement asymmetries: AI is frequently acquired as a turnkey black box from vendors, with limited bargaining power for audit rights, documentation, and model updates—conditions ripe for opacity and dependency.

DESIGNING AFRICA'S INTELLIGENCE FUTURE: Essays on AI, Governance, and Public Value

UCD, oriented to front-end interactions, insufficiently addresses these structural risks. The right unit of design is the socio-technical system: legal bases for processing, data retention schedules, appeals processes, third-party oversight, and sunset clauses. UNESCO’s Recommendation on the Ethics of AI emphasizes the importance of impact assessment, stakeholder inclusion, and proportionality, while the OECD principles prioritize human-centred values and robustness. Translating these into African settings requires procurement templates that demand algorithmic registers, auditability, and redress guarantees, as well as regulator capacity and public reporting on performance parity across demographic groups.

Public Value-Centred Design reframes ‘smart’ from optimization to accountability. The design question is not ‘What features do users want?’ but ‘What guarantees do citizens deserve?’ Guarantees include the right to human review for significant decisions, direct notification when automated systems are used, accessible explanations in local languages, and affordable channels for contestation. Here, the metrics of success are reductions in wrongful denials, improvements in explainability scores, and demonstrable parity across groups—not the number of AI features shipped.

8. DIGITAL SKILLS AND STATE CAPABILITY: DESIGNING FOR INTERMEDIARIES AND LEARNING BUREAUCRACIES

Digital transformation is capability-hungry: product management, service design, cybersecurity, data engineering, AI ethics, and change management. Many African administrations suffer capacity traps: skilled staff are thinly spread; rigid pay scales impede recruitment; vendor dependence is chronic. UCD’s field research can illuminate skill bottlenecks but cannot substitute for institutional reforms.

Designing for intermediaries is essential. Street-level bureaucrats, e.g., clerks, enumerators, social workers, as well as community-based organizations, translate policy into practice. Training them as co-designers and stewards, and equipping them with tools to override or annotate automated decisions, expands the state’s capacity to handle edge cases. Intermediaries can host assisted-digital sessions, help citizens assemble documentation, and escalate systemic defects detected at the coalface. The principle is not to ‘remove the human’ but to put the right human in the loop with the proper authority.

Organizationally, Public Value-Centred Design aligns with Problem-Driven Iterative Adaptation: anchor teams around specific, measurable problems (e.g., reducing the number of weeks to register a business in Region X without raising fraud), create space for time-boxed experiments with clear authorizing sponsors, and iterate with evidence. Capability is built through doing: internal digital corps, fellowship programmes with universities, and secondments from the private sector can seed skills. Procurement must evolve from waterfall specifications to outcome-based contracts that incentivize inclusion, robustness, and knowledge transfer, rather than focusing on feature counts.

9. CASE COMPARISONS: STYLIZED NORTHERN AND AFRICAN APPLICATIONS

TO SHARPEN CONTRASTS, CONSIDER THREE PAIRED SKETCHES:

Identity and data exchange. Estonia built X-Road after establishing a universal ID and registries; audit logs provide citizens with visibility into data access, operationalizing contextual integrity. UCD’s role focused on service findability and consistency on stable rails. In several African states, by contrast, identity and address registries are still being built; duplicative or siloed IDs coexist; and redress mechanisms are nascent. Here, UCD must be subordinated to the more complex work of infrastructuring identity quality, enrolment equity, and lawful processing. The core ‘user story’ is not a slick login but the guarantee that a pastoralist can enrol, that her record will not vanish in a breach, and that she can see and contest data flows.

Service access models. The UK’s GDS transitioned services to the web, incorporating a strong design system and inclusion guidelines. Assisted-digital support was developed through libraries and charities. In Kenya and Rwanda, hybrid one-stop centres (Huduma, Irembo) function as institutional intermediaries, recognizing the limits of pure digital self-service. The UCD lesson is not that ‘mobile-first’ always wins, but that channel strategy is a constitutional choice: who bears the burden of navigation and cost? Public Value-Centred Design mandates public investment to reduce citizens’ navigation costs when capability constraints are foreseeable—a principle of design solidarity.

Tax administration. Northern revenue agencies that adopted UCD improved forms and guidance while retaining robust audit and appeal structures. In African contexts, digitization of filing and payments has reduced leakages and friction for formal firms but sometimes widened gaps for informal traders who lack literacy or bank access, shifting them towards intermediated compliance that can be exploitative. Public Value-Centred Design would benchmark success not only by online filing rates but also by equity of burden: Are new systems raising compliance costs for the least capable? Are grievance channels functioning across languages and regions?

10. A NORMATIVE ALTERNATIVE: PUBLIC-VALUE-CENTRED DESIGN (PVCD)

If UCD is necessary but insufficient, what should replace it as the master philosophy? Public Value-Centred Design integrates Ubuntu, capabilities, and Problem-Driven Iterative Adaptation into a practicable doctrine for public digital transformation in Africa. It rests on five commitments.

From needs to guarantees. Shift the design question from ‘What do users need?’ to ‘What guarantees do citizens deserve?’ Guarantees include accessibility across languages and channels, explainability and redress for automated decisions, privacy by contextual integrity, and security commensurate with risk.

From satisfaction to legitimacy. Elevate metrics such as perceived fairness, trust, and realization of rights alongside usability. Embed grievance redress and independent oversight into service journeys; publish service-level objectives for response and appeal times; and practice radical transparency regarding outages and incidents.

From prototypes to authorizing environments. Design iteration within legal and budgetary authorization using Problem-Driven Iterative Adaptation. Treat legislation, standards, and procurement reform as part of the design space, not constraints to be worked around.

From individuals to intermediated ecologies. Resource-assisted digital channels, mobile kiosks, and community organizations as first-class parts of the service ecosystem. Train intermediaries in cybersecurity hygiene and data minimization; provide them with dashboards to surface systemic issues.

From features to infrastructure: Prioritize identity quality, data exchange governance, logging, and cybersecurity as foundational. Use threat modelling and red-teaming as routine design activities; publish algorithmic registers and perform impact assessments for high-risk AI.

Public Value-Centred Design does not discard UCD methods; it reorders them in accordance with constitutional ends. Empathy is an instrument, not a telos.

DESIGNING AFRICA'S INTELLIGENCE FUTURE: Essays on AI, Governance, and Public Value

11. OPERATIONALIZING PUBLIC VALUE-CENTRED DESIGN: A DECISION FRAMEWORK

Public Value-Centred Design can be operationalized through a ‘Public Value-Risk Design’ (PVRD) framework, a simple yet disciplined decision-making process for programmes.

Step 1: Diagnose the public value objective and risk surface. Define the collective aim (e.g., expand birth registration coverage by 20% in rural districts) and enumerate rights-affecting risks (exclusion, surveillance creep, breach). Map contextual conversion factors like literacy, connectivity and language diversity that condition capabilities.

Step 2: Identify stakeholders and intermediaries. Go beyond ‘end users’ to include street-level bureaucrats, community organizations, vendors, data protection authorities, and likely adversaries. Apply Arnstein’s ladder to design participatory processes that move beyond tokenism.

Step 3: Choose architectures and guarantees. Decide on federated vs centralized registries, data minimization practices, and logging/consent models. Specify guarantees: maximum decision turnaround times, human-review thresholds, and multilingual notices. Select security controls based on threat modelling.

Step 4: Iterate within authorization. Use Problem-Driven Iterative Adaptation to run time-boxed experiments with clear political sponsors and budgeted risk. Build learning loops: publish weekly metrics on inclusion, redress, and security incidents; hold public ‘show the thing’ sessions with civil society.

Step 5: Institutionalize accountability. Enact or update legal bases, empower regulators, establish independent audit panels for algorithms, and create public registers of automated systems. Tie vendor payments to outcome-based milestones, including metrics related to inclusion and security.

This framework welcomes UCD activities such as research, prototyping, and usability testing at Steps 2 and 4, but requires that they serve guarantees established in Step 3 and accountability in Step 5.

Step 6: Measure what matters. Replace vanity metrics—such as page views and completion times—with public-value indicators: equity of access across gender, language, disability, and geography; procedural justice (were reasons given? were appeals accessible?); and resilience (mean time to recovery after incidents). Track ‘capability conversion’ by measuring whether services translate into tangible outcomes such as birth certificates used to enrol children, tax registrations leading to access to finance, rather than counting tokens of interaction. Incorporate feminist and decolonial data insights to interrogate who is missing from datasets and who bears the burden of data production. Finally, make legitimacy an explicit KPI by surveying perceived fairness and trust before and after interventions. What gets measured gets managed; measure the right things.

12 ANTICIPATING OBJECTIONS

Three objections merit a response.

Objection 1:

Public Value-Centred Design will slow innovation.

Reply:

In the public sector, moving fast and breaking things breaks rights. Public Value-Centred Design accelerates what matters by avoiding rework from security failures and legitimacy crises. Problem-Driven Iterative Adaptation’s small bets are faster than stalled big-bang projects.

Objection 2:

Public Value-Centred Design is paternalistic; users are aware of their own needs.

Reply:

Public Value-Centred Design does not deny agency; it broadens the moral community and recognizes asymmetries. In contexts of low literacy and power imbalance, the state has a duty to provide guarantees that expand capabilities and prevent harm.

Objection 3:

Public Value-Centred Design is too abstract for overstretched governments.

Reply:

The framework’s steps translate into concrete checklists (algorithm registers, DPIAs, grievance flows, threat models) that can be templated and shared across agencies. International bodies already provide scaffolds; the task is localization and enforcement.

13. TOWARDS A PLURIVERSAL DIGITAL STATE

A call to design for the pluriverse invites African states to resist a one-size-fits-all digital modernity. Ubuntu grounds a civic imaginary in which services cultivate solidarity and dignity rather than merely streamline transactions. The capability approach insists that what matters is what people can actually do with services, not whether they clicked a button. Problem-Driven Iterative Adaptation reminds us that capability grows through practice under authorization, not through imported blueprints.

Comparisons with Northern successes should be humbling and clarifying. Estonia’s sleekness is inseparable from decades of institution-building; the UK’s design leadership rides on statutory rights and regulators; U.S. rescue teams succeed when they can commandeer authority to cut through procurement knots. Africa’s significant digital breakthroughs from mobile money ecosystems, hybrid one-stop centres, and community health information systems emerged when design respected intermediaries and worked within political realities rather than against them.

Regional cooperation can accelerate this shift. Continental and regional instruments—the African Union’s Digital Transformation Strategy 2020–2030, the Data Policy Framework, and the Convention on Cyber Security and Personal Data Protection (Malabo Convention)—offer scaffolds for interoperable identity, data sharing, and safeguards that transcend national capacity constraints. Pooling expertise through regional ‘public interest technology’ hubs, shared security operations centres, and cross-border sandboxes can reduce vendor lock-in and strengthen bargaining power for algorithmic audit rights and open standards.

CONCLUSION:
FROM DELIGHT TO DIGNITY

UCD gave the public sector powerful tools: ethnographic humility, iterative building, and a language for centring human experience. However, its underlying anthropology and metrics are misaligned with the moral project of the state, especially in African contexts where plural publics, infrastructural work, and legitimacy deficits define the terrain. A provocative yet practical reframing—Public-Value-Centred Design—places guarantee over needs, legitimacy over satisfaction, authorization over skunkworks, intermediaries over atomized end-users, and infrastructure over features. It invites designers to treat law, procurement, and security as design materials; to stage participation as power-sharing, not theatre; to measure capability expansion rather than clicks; and to treat algorithms as accountable public servants, not inscrutable oracles.

If there is a single design maxim to carry forward, let it be this: treat the rule of law, the right to explanation, and the capability to access as first-class interface components. When citizens see their dignity reflected not only in the smoothness of clicks but in the fairness of outcomes, the clarity of reasons, the safety of their data, and the availability of help when systems fail, Africa’s digital state will have moved beyond empathy to justice. That destination is not anti-user; it is truer to the social contract. In the long run, dignity scales; gimmicks do not.

Design for justice is harder, slower, and ultimately the only way democratic digital states endure. For everyone, everywhere.

DATA AS DESTINY:

Unlocking Efficiency and Trust in Public Service

NILOFAR SAIYED

INTRODUCTION

Imagine a scenario where one ministry provides 100 percent toilet access to underprivileged families, yet these facilities remain unused because a different ministry, responsible for water supply, has not extended the necessary connections. Or consider a situation where one ministry possesses all the land-use data, while another struggles to levy appropriate user charges for basic services due to a critical absence of corresponding land measurement information.

These are not isolated incidents; they highlight a persistent issue: without a modern approach to data, progress on vital schemes will always be slow. The question, then, is what exactly are the hidden costs of our flawed data ecosystems?

1. FLAWED POLICY AND MISGUIDED BUDGETS

Without timely, accurate, and comprehensive data, leaders are forced to make decisions based on guesswork, leading to policies that are not aligned with on-the-ground realities. The poor quality of existing data often compounds this issue, as analysts spend as much as two-thirds of their time cleaning and organizing bad data rather than on analysis. This dramatically slows down processes and hampers innovation. The Organisation for Economic Co-operation and Development (OECD) has noted that data used for policy-making is often unreliable, leading to a risk of ‘policy failures’ and the wasteful duplication of efforts. Without a clear evidence base, budgets are allocated based on historical trends or political expediency rather than actual needs. This can result in essential funds being directed to projects in areas that do not require them, while truly vulnerable populations are left unserved.

The World Bank emphasizes that data is critical for providing ‘timelier and finer-scale assessments of programs,’ and without it, governments risk significant misallocations of public funds. This problem is widespread. In Ghana, for instance, a substantial lack of reliable data on smallholder farmers, including crop yields, soil types, and market access, has hindered the government's ability to design effective agricultural support programmes, directly impacting food security and economic growth. Similarly, a report on India's finances found large variations in underreporting across state governments, with actual spending on food subsidies being far higher than reported due to data gaps. These are not just accounting errors; they represent a fundamental inability to accurately track public spending and ensure it is making the intended impact.

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

2. THE COST OF INSTITUTIONAL FRAGMENTATION: REDUNDANCY AND DELAYS

The inefficiency of government is often a direct result of ministries operating in silos, viewing their data as a private asset. This institutional reluctance to share information with other departments creates significant delays and hinders collaborative efforts. The World Bank's own research finds that this lack of data sharing is a major hurdle to effective governance. A key reason for this is that ministries often lack a clear data classification protocol, which is often used as a convenient excuse to hoard data, as ministries can claim to be protecting sensitive information without a standardized framework to guide them. This leads to redundant data collection and prevents the creation of a unified picture of national needs, which is essential for effective, cross-sectoral planning. Tata Consultancy Services reported that this lack of coordination was a key factor in why the COVID-19 pandemic response was slowed, as critical data on disease spread could not be shared rapidly across sectors.

3. THE ACCOUNTABILITY GAP: THE COST OF UNMONITORED SCHEMES

Leaders who are committed to bringing about change find their efforts hindered by an inability to effectively monitor the implementation of well-designed schemes. The data, often lacking accountability and a proper format, is not readily available on digital dashboards for a real-time view of progress. This makes it incredibly difficult to hold officials accountable for their performance. As a result, transparency becomes a major issue, hindering NGOs, the public, and civil society from participating in oversight and ensuring that services reach the intended beneficiaries.

The Harvard Data Science Review states that without high-quality data that is available and comparable over time, sustainable development is ‘doomed to falter’. This is a reality in many parts of the world, including Sub-Saharan Africa, where

a World Bank report highlighted that over 40 percent of births in more than half of the region's countries are not registered, creating a data gap that jeopardizes children's access to essential services. The lack of reliable data directly impacts public trust, as citizens are less likely to trust institutions that are not transparent about their operations.

The scale of these deficiencies presents a direct challenge to leadership. But what if we could turn this challenge into a profound opportunity? A new era of data-driven governance is not merely an option; it is an imperative for sustainable progress.

THE IMPERATIVE OF DATA-DRIVEN GOVERNANCE

The systemic issues that slow progress and erode public trust are not a permanent condition; they are a direct challenge to leadership. A new era of data-driven governance offers a strategic path forward, a way to move beyond the inherited inefficiencies of the past and build a lasting legacy of effective, citizen-focused government. This is a choice that will determine not only the speed of national progress but the very reputation of a leader committed to genuine transformation.

Leaders can now gain a clear, validated view of their impact in real-time, moving away from fragmented reports and guesswork. A robust digital ecosystem, powered by AI-enabled dashboards, provides a live and comprehensive view of a programme's implementation. This gives leadership a powerful tool to:

Proactively steer national initiatives: Rather than waiting for a yearly report, leaders can monitor key performance indicators of welfare schemes instantly, identify bottlenecks as they emerge, and ensure ministries are collaborating effectively. This empowers them to act decisively and visibly, demonstrating a hands-on commitment to results.

Build policies on validated truths: A digital ecosystem can automatically validate data from multiple sources, reducing the risk of false reporting. This ensures

that the information guiding a leader's decisions is accurate and authentic, providing a foundation of truth that is difficult to challenge.

Judicious use of national resources:
By implementing a unified data-sharing platform, a government can apply the ‘once-only’ principle, collecting data just once and sharing it securely across authorized departments. This approach, by its very nature, eliminates the need for redundant data collection, which saves immense time and resources.

Empowering a more capable government:
Any public department that relies on data, which includes nearly all of them, can operate with unprecedented efficiency and efficacy. For instance, a public health agency can more accurately forecast disease outbreaks by securely accessing anonymized data on patient visits from hospitals across the country. The UN Department of Economic and Social Affairs highlights that such data integration is critical for generating evidence that can inform both short-term responses and long-term policies, thus ensuring judicious use of the nation’s resources.

Leaders who champion data-driven governance leave behind a legacy that is durable and deeply trusted. These reforms are not just about technology; they are about fundamentally reforming a centuries-old system for the better.

Rebuilding public trust:
A transparent, digital system provides a powerful tool for accountability. It makes it easier to measure performance, identify failures, and hold officials responsible. This transparency, when institutionalized, helps to rebuild the public trust that is so often eroded by inefficiency and corruption.

Holistic development begins with data:
A digital data governance system is the foundational step for any kind of holistic national development. It creates the stable, transparent, and efficient ecosystem needed for progress in healthcare, education, economic growth, and citizen welfare. Leaders who invest in this fundamental infrastructure demonstrate a commitment to building a more resilient, adaptive, and accountable future for all.

The promise of data-driven governance is not an untested theory. It is being put into practice by leaders around the world who are reshaping public service delivery. Their experiences offer invaluable lessons and a clear roadmap for others to follow.

**ACCELERATING TRANSFORMATION:
GLOBAL BEST PRACTICES IN DATA-DRIVEN
GOVERNANCE**

The challenge of leveraging data for national development is a global one, and leaders around the world are showing us how it is done. The most effective strategies are not about complex technology, but about clear vision and deliberate choices. The lessons from these experiences offer a playbook for any nation ready to own its digital future.

1. FOUNDATIONAL VISION AND STRATEGIC PLANNING:
Successful data-driven governance often begins with a decisive, high-level commitment. For instance, Saudi Arabia's leadership positioned data and AI as a central pillar of its economic transformation agenda with a clear and ambitious National Strategy for Data & AI (NSDAI). This is a vital first step for any government seeking to drive large-scale reform by treating data not as a secondary concern, but as a primary resource for national strategy.

2. BUILDING TRUST THROUGH SECURE GOVERNANCE:
With a clear vision in place, the next step is to establish foundational governance. The UK government, for example, reformed its data classification system to a simplified three-tiered model, which allowed for the secure adoption of public cloud services and unlocked a significant amount of public data for government use. As a result, over 90 percent of UK government data was classified as ‘Official’, allowing it to be stored and processed securely in the public cloud. Similarly, Singapore’s ‘Cloud First’ policy and Malaysia's Data Sharing Bill are powerful examples of how clear legislative frameworks can break down institutional silos and drive efficiency. These cases highlight the critical importance of a clear data classification protocol as a foundational step for modern, efficient governance.

**DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value**

3. ENHANCING EFFICIENCY WITH INTEROPERABILITY:
Once the groundwork is laid, leaders can focus on building systems that deliver real-world results. Estonia, a pioneer in e-governance, developed X-Road, a secure data exchange layer that operates on the powerful ‘once-only’ principle. This means that citizens only have to provide their data to the government once; it can then be securely shared internally between different agencies. The system saves the country an estimated 1,345 years of working time annually by eliminating redundant data collection and paperwork. Similarly, India’s DARPAN dashboard provides leaders with a real-time, quantifiable view of welfare projects, replacing guesswork with hard facts and empowering leaders to hold officials accountable instantly.

4. DELIVERING TANGIBLE BENEFITS TO CITIZENS:
Ultimately, these initiatives culminate in tangible benefits for citizens. The Huduma Centres in Kenya brought over 50 government services under one roof, serving millions of citizens annually. This is a powerful model for citizen-centric service delivery. In addition, India's Aadhaar digital identity system shows how a single, foundational data tool can fight corruption and drive financial inclusion on a massive scale. By linking welfare benefits directly to citizens’ biometric data through the Direct Benefit Transfer (DBT) system, the government has saved billions by eliminating fraudulent beneficiaries and ensuring that support reaches those who need it most.

The journey to data-driven governance is not a distant promise. It is a strategic choice available to any leader ready to make it and claim their data as destiny. The following section outlines a strategic roadmap for institutional reform.

**A STRATEGIC FRAMEWORK FOR
INSTITUTIONAL REFORM**

**STRENGTHENING INSTITUTIONAL AND LEGAL
FOUNDATIONS FOR DATA-DRIVEN GOVERNANCE:**

1. ESTABLISH CENTRALIZED LEADERSHIP FOR DIGITAL DATA GOVERNANCE:
Institutional reform begins with decisive leadership. To overcome traditional fragmentation, a single, powerful entity such as a Ministry of Information Technology could be given the mandate and authority to lead this transformation. This central body would serve as the coordinating authority for all departmental initiatives, ensuring a unified approach and preventing the duplication of efforts. Its purpose is to formulate a comprehensive governance framework that provides the essential rules of engagement for leaders, civil servants, and technical teams to ensure seamless monitoring and effective implementation of welfare schemes in a way that is coordinated with manpower resources of different ministries and policy changes when required.

2. DEVELOP A ROBUST LEGAL AND REGULATORY FRAMEWORK:
A foundational recommendation for any digital data ecosystem is to establish a strong legal and regulatory framework. This involves moving beyond a simple data policy to a comprehensive data protection law that provides a clear legal basis for data sharing while safeguarding citizens' privacy rights. This framework could also define the roles and responsibilities of data controllers, set penalties for misuse, and establish an independent authority to provide oversight.

INFRASTRUCTURE AND DATA INTEGRITY:

1. DEFINE DATA CLASSIFICATION AND SHARING PROTOCOLS:

Before any data can be shared, a comprehensive data classification protocol should be in place. This process, as demonstrated by the UK’s reform to a three-tiered system, provides a common language for all government bodies. The framework should define how different data types are to be handled, where they are stored, how they are securely shared, and for how long they are retained. By formulating these protocols, the central body empowers ministry leaders to coordinate effectively and ensures a holistic approach to citizen welfare, preventing the inefficient duplication of efforts and resources.

2. ADOPT A STRATEGIC HYBRID CLOUD INFRASTRUCTURE:

A leader’s investment in infrastructure should be strategic and fiscally sound, rather than an expensive, upfront commitment to new data centres. A hybrid cloud strategy offers a balanced approach. This allows less-sensitive public data to be stored cost-effectively on commercial cloud services, while highly confidential data is kept on secure, government-owned servers. This approach aligns with the practices of nations like Singapore and the UK, which leverage public clouds for scalability while maintaining sovereign control over critical information.

3. BUILD A VALID DATA PIPELINE FROM THE GROUND UP:

The integrity of the entire data ecosystem is only as strong as the data that feeds it. To ensure data validity and coherence, a comprehensive strategy for data feeding and verification should be implemented. The solution is to equip field staff with mobile applications for data collection. These apps could have built-in validation rules, GPS tagging, and time-stamping to verify the location and authenticity of data entry. For any progress or changes in a situation on the ground, a note explaining the reason for the change should be entered, providing a clear audit trail and strengthening accountability.

4. CREATE AN INCLUSIVE AND SMART DATA FEEDBACK LOOP: A key innovation for cohesive governance is to develop a single, standardized questionnaire for welfare-related data. This questionnaire would collect information relevant to multiple ministries simultaneously. To ensure no citizen is left behind, especially those who are not digitally literate, an Interactive Voice Response (IVR) system can be used to gather updates. AI models can then analyse this IVR data, converting it into structured formats that are seamlessly fed into the databases of all relevant ministries. This strategy ensures the data of all welfare schemes is valid and updated periodically, with no discrepancies between different ministries' records.

CULTURE AND ACCOUNTABILITY

1. CULTIVATE A CULTURE OF DATA-DRIVEN ACCOUNTABILITY:

Technology is only as effective as the people who use it. Systemic reform requires a commitment to training and fostering a new culture of collaboration and accountability. A comprehensive training programme is required for government employees, from field officers who gather data to top-level bureaucrats who oversee the systems. This involves creating a culture where data is the standard for debate and a basis for accountability. The National Digital Council, in coordination with the lead ministry, could make the use of real-time dashboards and validated data a core part of performance management.

2. BUILD A COLLABORATIVE MODEL WITH EXPERTS AND THE PRIVATE SECTOR:

The success of this framework hinges on a ‘whole-of-society’ approach. The central body could actively involve NGOs for insights on data collection formats that are practical for field staff and engage with startups and local IT companies to develop agile solutions. Furthermore, inviting inputs from global technology leaders and top companies is crucial to introduce world-class expertise and ensure the capacity of national companies is built to sustain the ecosystem's long-term success. It is also strategic to send key civil servants and leaders abroad to study best practices firsthand and bring those insights back to inform policy. This is particularly critical when negotiating with external technology providers.

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

As highlighted in Sitong’s essay, Smart Cities, Smart Contracts: Negotiating AI Partnerships on African Terms, a clear understanding of the negotiation process is essential to avoid vendor lock-in and ensure that data is used to serve the nation rather than undermine it.

PARTNERSHIPS AND LONG-TERM SUSTAINABILITY

1. PROMOTE PUBLIC-PRIVATE PARTNERSHIPS (PPPS) FOR DIGITAL INFRASTRUCTURE:

For countries with limited resources, PPPs offer a powerful mechanism to accelerate digital transformation. Governments can partner with the private sector to develop and maintain digital infrastructure, such as data centres and broadband networks, that would otherwise be too costly to build and operate alone. This approach allows the government to focus on policy and oversight while leveraging private capital and expertise to bridge the infrastructure gap.

2. ENSURE LONG-TERM ECONOMIC SUSTAINABILITY AND ENVIRONMENTAL RESPONSIBILITY:

The final recommendation is to view this digital ecosystem as a long-term investment that is both economically and environmentally sustainable. The cost of establishing and maintaining this system could be justified by the significant benefits it generates, such as saving an amount equivalent to 0.8-1.1 percent of GDP annually by reducing fraud and administrative costs. Furthermore, the government should proactively address the ecological footprint of its digital infrastructure by using energy-efficient hardware and sustainable cooling technologies, ensuring that national development is aligned with environmental goals.

The recommendations presented here lay out a comprehensive strategy for building a modern, citizen-centric data ecosystem. This path requires a sustained commitment from leadership, and its ultimate success will be defined by the legacy it leaves for the country and its citizens.

**CONCLUSION:
A LEGACY OF TRUST AND TRANSFORMATION**

The decision to invest in a national data ecosystem is not merely a technical upgrade but a choice to redefine governance itself. For African nations, it offers a path to move beyond fragmentation and build systems grounded in evidence, transparency, and trust. When data is shared responsibly across institutions, decisions become timely, resources are used wisely, and citizens experience the state as responsive and fair. This transformation is ultimately about leadership—about choosing to treat data not as an administrative tool but as a catalyst for accountability and inclusive growth. By owning their data, nations can own their destiny, shaping a future where technology serves governance and governance, in turn, serves the people.

THE RAILS OF INCLUSION:

Building Africa’s Digital Foundations for AI

ARFAN RAMADHAN

Imagine the next pandemic, the next flood, the next price shock. A government will press a button—and for some households, help will arrive in minutes. For others, nothing. Not because leaders didn’t care, but because the digital rails were never built to carry everyone.

This is the uncomfortable truth: in much of Africa, millions of citizens remain unseen by the very systems now being asked to power an AI-enabled future. Refugees without paperwork, women in rural villages, families in informal settlements—they are absent from official records, absent from payment rails, absent from the data pipelines that drive decisions. In an era where artificial intelligence feeds on data, invisibility is not just an injustice. It is a design flaw.

And yet, the story is not one of despair, but of opportunity. Across the continent, we see glimpses of what is possible when inclusive rails exist. In Togo, during COVID-19, emergency aid reached thousands through mobile money, guided by machine learning models that worked only because payments and identity were aligned. In India, a consent-based system now allows citizens to control how their personal data moves, proving that trust and innovation can grow together. In Estonia, citizens can see an audit trail of every data query, a daily reminder that sovereignty is not only national but personal.

What Africa faces today is a choice: to inherit AI on someone else’s terms, or to build its own foundations for a digital state that is inclusive by default. These foundations are not glamorous, but they are indispensable. I call them the Rails of Inclusion—three essential commitments that ensure no one is left behind in the exponential age:

- Trusted digital identity:**
So that every person is visible, counted, and able to claim their rights.
- Interoperable payment systems:**
So that value can move instantly and fairly to those who need it.
- Safe data exchange:**
So that information flows with consent, transparency, and accountability.

Together, these elements form the precondition for any equitable use of AI in Africa. Without them, automation will reproduce yesterday’s inequalities at digital speed. With them, AI becomes a lever for precision inclusion—ensuring that the future does not arrive for some, but for all. This essay explores how Africa can lay down these rails of inclusion. It looks at the design choices that matter, the pitfalls to avoid, and the opportunities to seize. Above all, it argues that building these foundations is not a technical task alone, but a moral one: the work of shaping a digital future where every citizen is seen, trusted, and empowered.

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

**THE PROBLEM:
INVISIBILITY BY DESIGN**

The paradox of our time is that governments are embracing artificial intelligence faster than they are fixing the foundations on which it stands. The result is an ‘acceleration gap’: the tools of tomorrow are running ahead on the tracks of yesterday. And when those tracks are incomplete, people are left behind.

The first gap is identity. Across the world, some 850 million people still lack official ID—many of them in Sub-Saharan Africa. Without recognition on paper, they are invisible to systems that mediate life itself: opening a bank account, visiting a clinic, registering a SIM card. In an AI-mediated state, this gap becomes even sharper. Models trained on partial datasets will misallocate resources, not because the algorithm is malicious, but because the citizen was never recorded in the first place. Exclusion is not a glitch; it is baked into the system.

The second gap is representation. African ancestry accounts for just over one percent of global genomics datasets. Medical AI trained on such skewed samples produces brittle models that travel poorly across borders. A diagnostic tool that works in Boston may fail in Bamako. Without local data and governance, the promise of precision medicine risks becoming the practice of precision exclusion.

The third gap is algorithmic disparity. Studies show that commercial facial analysis systems make far more errors for darker-skinned women than for lighter-skinned men. When these systems are deployed in e-KYC, border control, or public safety, the errors do not fall evenly. They concentrate precisely on those already at risk—women, minorities, the displaced—multiplying existing vulnerabilities at digital speed.

The fourth gap is trust in data exchange. Ministries often hoard information out of fear, or share it in opaque and ad-hoc ways. The result is a lose-lose: either innovation stalls because no one can access the data, or rights erode because data flows without safeguards. Yet mature counter-models already exist—from Estonia’s transparent audit trails to India’s consent managers. These prove that it is possible to share data in ways that are both innovative and respectful, both fast and fair.

Finally, there is the gap of payments. AI cannot target what it cannot pay. Interoperable, auditable payment systems are the delivery layer of inclusion. Where rails are fragmented or always-online only, cash transfers fail to reach the poorest precisely when they are most needed. During crises, this is not an abstract failure—it is food not delivered, medicine not purchased, dignity deferred.

Africa has already seen both sides of the story. In Kenya, the National Integrated Identity Management System (NIIMS), known as Huduma Namba, was halted twice by the courts due to the absence of an adequate data protection framework, raising risks of privacy breaches, exclusion of marginalized groups (e.g., through biometric failures), and misuse of sensitive biometrics. This project taught us a vital lesson: inclusivity and safeguards must be built in from the start, so that progress does not stall. The court interventions were not merely setbacks, but reminders that legitimacy and trust are as important as technical design. In Togo, by contrast, mobile money rails enabled emergency cash transfers during COVID-19, demonstrating the power of even partial alignment between identity, payments, and data.

The lesson is clear. Without trusted identity, inclusive payments, and safe data exchange, artificial intelligence will amplify invisibility. With them, AI can become a lever of precision inclusion, ensuring that the next crisis does not repeat the inequities of the past but begins to close them.

**THE FIRST ESSENTIAL ELEMENT:
DIGITAL IDENTITY—EVERYONE VISIBLE,
EVERYONE COUNTED**

Identity is the ticket to the state. It is the key that opens schools, clinics, banks, and welfare programmes. Without it, millions remain unseen—their lives unfolding outside the view of systems that are supposed to serve them. In an age where artificial intelligence feeds on data, invisibility is not just a social injustice. It is a structural weakness.

Legal identity is recognition before the law; a document is proof of that recognition. Digital identity adds the rails that let people prove who they are, reliably and repeatedly, across both online and offline channels. When governments build these systems inclusively, they ensure that no one is left behind. When they neglect safeguards, they risk hardening old gaps.

Yet, leaders cannot assume that enrolment requirements like birth certificates or passports are universally available. In the Sumatera floods of late November 2025 in Indonesia, this fragility became painfully visible. Survivors who had lost everything were barred from receiving food aid because the aid distribution centre required a personal or family identity card before distributing supplies. Some who had escaped the rising waters came close to starvation—not because aid was unavailable, but because they could not prove who they were. The system demanded documents; the people needed relief.

This is why governments must take the state to where people live—markets, schools, refugee camps, disaster areas—making first-time issuance free of charge and allowing trusted community witnesses to attest where documents are missing. Visibility must begin at the edge.

Centralization also brings both promise and peril. When governments merge registries, they can reduce duplication and improve accuracy. But if they ignore exclusion risks, they simply replicate injustice at a larger scale. A central source of truth is not the same as a single source of justice. To avoid this trap, leaders must embed safeguards into the very architecture: collecting only the minimum data, prohibiting sensitive attributes such as DNA or GPS, and publishing audit trails so citizens can see who accessed their records and why.

Connectivity adds another challenge. Identity that works only when the network does is not public infrastructure. Citizens in rural villages cannot wait for coverage to stabilize before being recognized by the state. Leaders must design systems with resilience: offering biometrics where feasible, but also PINs, USSD codes, and offline verification so that no one is turned away because a fingerprint scanner failed or the signal dropped.

And because no system is flawless, governments must also guarantee redress. Mistakes will happen—but dignity is measured not by perfection, rather by fairness when errors occur. Appeals must be simple, services must continue during investigations, and errors must be corrected transparently.

As Nilofar reminds us in her essay *Data as Destiny: Unlocking Efficiency and Trust in Public Service*, progress falters when data systems are fragmented or unaccountable. A universal identity layer is one way to ensure the state sees every citizen, not just a partial picture. It is the bedrock for data ecosystems that are whole, trustworthy, and inclusive.

If governments succeed, the difference is tangible. A mother in a rural village can verify her identity on a basic phone, even when offline. A displaced family can enrol without being turned away for lack of papers. Citizens can see who accessed their records and hold institutions accountable for misuse. Identity becomes not a hurdle but a gateway—one that opens safely to many doors, without demanding re-enrolment at every step.

Building digital identity as public infrastructure is not about surveillance or efficiency alone. It is about recognition. It is about ensuring that in moments of crisis—a flood, a pandemic, a food shortage—leaders know every citizen is visible, and every life counts. When identity is inclusive at the edge and rights-preserving at the core, it becomes the firm ground on which AI can stand.

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

**THE SECOND ESSENTIAL ELEMENT:
SECURE PAYMENTS—MONEY THAT MOVES
FOR EVERYONE**

If identity tells us who is eligible, payments decide when and how value reaches them. In the next flood, drought, or price shock, governments will not mail cheques. They will press ‘send’—and millions of households will feel it, or not, within hours. A resilient payment rail is therefore not a luxury. It is the nervous system of inclusion.

When rails are fragmented or fragile, people are stranded. If each bank or wallet operates in its own bubble, or if every transfer depends on an internet connection, households are cut off precisely when help is most urgent. Governments must treat payment rails as a public utility—an interoperable backbone open to banks, mobile-money providers, and fintechs alike. Just as highways carry all vehicles under the same rules, shared standards should allow money to flow seamlessly across providers, even when offline.

Payments must also be accessible and trusted. Too often, opening an account requires documents that many citizens do not have. Others hesitate because every transaction is tied automatically to national IDs, raising fears of surveillance. Leaders can overcome this by enabling tiered access: light checks for small, low-risk accounts and stronger ones for larger sums. They must ensure that data linkages are always consent-based, not automatic. Citizens will only use a system they trust—and AI can only optimize flows that people choose to enter.

Equally important is fairness and cost. High fees and opaque charges quietly erode the value of every transfer. Weak dispute mechanisms leave citizens powerless. Leaders must set transparent rules for pricing, guarantee fast and fair recourse for errors, and subsidize public-interest uses such as emergency transfers. If agents and merchants cannot afford to participate, the rails will collapse.

If citizens believe the system is unfair, they will opt out. Both outcomes weaken the very foundation on which AI-driven targeting depends.

The promise of resilience is already visible. In Niger, households receiving mobile-money transfers improved their diets and saved hours of travel compared to in-person cash collection. In Ghana, interoperability allowed citizens to send money across networks instantly, fuelling both commerce and trust. And regionally, the Pan-African Payment and Settlement System is opening doors for instant cross-border transfers, laying the groundwork for AI models that can forecast shocks and optimize trade across the continent.

Sitong’s essay on *Smart Cities, Smart Contracts: Negotiating AI Partnerships on African Terms* reminds us that power is negotiated, not given. The same principle applies to payments: governments must negotiate rules that prevent vendor lock-in and guarantee interoperability, so no single provider can hold citizens’ lifelines hostage.

When payments work, the impact is tangible. A farmer in a remote ward receives emergency support within hours, even when the signal is weak, because offline fallbacks are built in. A young trader moves money across borders without prohibitive costs or delays. Citizens see clear fee schedules and know they can challenge errors without fear. Trust grows because fairness is visible.

Payments are not just pipes for money. They are the delivery mechanism of dignity. When governments build them to be interoperable, affordable, private by design, and resilient against disruption, they create the conditions for AI to serve rather than exclude. In times of crisis, money that moves quickly and fairly is more than a transaction—it is a lifeline.

**THE THIRD ESSENTIAL ELEMENT:
SAFE DATA EXCHANGE—TRUSTWORTHY FLOWS,
TRUSTED OUTCOMES**

If identity tells us who people are, and payments deliver what they are entitled to, then data exchange governs how systems talk to each other. For AI-enabled public services, safe data exchange is the oxygen: without it, models cannot be trained, tested, or deployed responsibly. But the same flows that make innovation possible can also entrench surveillance or spark distrust. The question is not whether data should move, but how—under whose authority, with what safeguards, and for whose benefit.

When data sharing is opaque, citizens feel excluded from the bargain. If people cannot see who accessed their information, or for what purpose, trust erodes. Governments must therefore log every data exchange, and citizens must be able to view those logs. Transparency is not a burden; it is the precondition of legitimacy.

When governments are tempted into over-collection, risk multiplies. Bigger databases do not always mean better models. Collecting DNA or GPS where it is unnecessary only heightens fears of intrusion. Leaders must learn to practise restraint: collecting the minimum attributes needed and nothing more. High-quality, relevant data serves AI better than bulk surveillance ever could.

Fragmentation creates another barrier. When ministries adopt incompatible systems, integration stalls and AI outcomes are skewed. Leaders must insist on interoperability built with safeguards: different agencies should be able to communicate, but only with modular access. The answer is not to replicate full dossiers, but to enable specific, auditable exchanges.

And because mistakes are inevitable, governments must guarantee redress and accountability. Citizens must have simple channels to appeal, independent regulators to oversee enforcement, and visible sanctions for misuse. Dignity is preserved when errors are correctable, not hidden.

Here, the lessons of other essays resonate strongly. As Nilofar argues, data is the lifeblood of sustainable development—without integrity and accountability, progress falters. And Sitong reminds us that sovereignty is negotiated in every contract. If African governments sign away control of data flows, they risk building dependence rather than resilience. Together, these insights point to the same truth: data rails must be both trusted and sovereign if they are to serve the public good.

Examples from elsewhere show this is possible. Estonia’s X-Road gives citizens a transparent audit trail for every data query, proving that accountability can be embedded in infrastructure itself. India’s Data Empowerment and Protection Architecture allows individuals to authorize data flows through consent managers, demonstrating that innovation and choice can coexist. The African Union’s own Data Policy Framework offers a continental blueprint, but the real test is national implementation.

What success looks like is clear. A patient can see who accessed her medical records and challenge misuse. A farmer’s information moves securely between ministries without exposing his entire file. Citizens know that consent is not a one-time formality but a revocable choice. Ministries exchange data lawfully, under shared standards, with independent oversight at the table. Trust rises, participation expands, and AI becomes more representative and reliable.

Safe data exchange is not only about pipes and protocols. It is about power. It is about whether citizens believe that digital government serves them, or spies on them. When governments design data rails to be transparent, consent-based, and auditable, they do more than enable AI. They build legitimacy, sovereignty, and trust—the very foundations of a digital state that is inclusive by default.

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

**THE RAILS OF INCLUSION—
A COHERENT FRAMEWORK**

Taken together, digital identity, secure payments, and safe data exchange form what I call the Rails of Inclusion. Each is powerful on its own, but their true strength lies in how they reinforce one another.

Identity without payments leaves recognition without delivery. A citizen may be visible in the registry but unable to receive a subsidy or transfer when crisis strikes. Payments without data exchange create islands of access, unable to scale beyond one-off schemes. Data exchange without identity risks feeding models with incomplete truths, or moving information without the consent of those it concerns.

Only when these three elements are aligned can artificial intelligence serve as a lever for inclusion rather than exclusion. Together, they create the conditions for precision delivery: governments can target aid quickly, citizens can claim entitlements reliably, and services can adapt dynamically to new shocks.

This is not theoretical. We have already seen glimpses of the whole system in practice. In Togo, mobile payments linked to basic identity records allowed emergency transfers to reach households during COVID-19. In Estonia, interoperable data exchange ensured that every query was transparent and accountable. In India, consent managers put citizens in control of how their data moved. These examples, drawn from diverse contexts, point towards the same truth: inclusive rails make AI both more effective and more just.

Building these rails does not make headlines. But it is the indispensable groundwork for a future where AI strengthens legitimacy, equity, and trust. Without them, governments risk building digital states for some, not all. With them, Africa can shape an AI-enabled future that is inclusive by default.

**CONCLUSION:
FROM INVISIBLE BY DESIGN TO INCLUSIVE BY DEFAULT**

The rush to adopt artificial intelligence across Africa is often framed as a race—a sprint to catch up with global leaders before it is too late. But the real contest is not about speed. It is about visibility. It is about whether every citizen is seen, trusted, and served by the digital state we are building. Once again, the Rails of Inclusion—trusted identity, secure payments, and safe data exchange—are not something spectacular. They do not carry the drama of frontier models or the spectacle of global competitions. But they are the foundations without which AI cannot serve fairly. They are the quiet architecture of legitimacy.

This work is also deeply political. It requires leaders to challenge entrenched interests that thrive on siloed data, to resist vendors who promise shortcuts at the cost of sovereignty, and to invest in infrastructure that is inclusive by design. It is not easy, but it is essential. As Nilofar shows, data without accountability slows progress. As Sitong reminds us, sovereignty without negotiation is an illusion. The Rails of Inclusion respond to both calls, weaving development and sovereignty into one project.

The choice is clear. African governments can inherit yesterday’s inequalities at digital speed, or they can build tomorrow’s legitimacy through inclusive rails. They can allow AI to magnify existing blind spots, or they can use it to close the gaps.

The future is not waiting. Floods, pandemics, and price shocks are already testing our systems. The question is whether the next ‘send’ will reach everyone, or only some. By investing in the Rails of Inclusion now, leaders can ensure that Africa’s AI future is Not Invisible by Design, but Inclusive by Default.

SMART CITIES, SMART CONTRACTS:

Negotiating AI Partnerships on African Terms

SITONG LYU

Sovereignty isn’t declared—it’s negotiated, and it’s negotiated in contracts. Africa’s AI future will be written in contract negotiations with external AI service providers. Every line shapes who owns the data, who carries the risk, and who captures the value. For leaders deciding whether a ‘smart’ system will strengthen sovereignty or strain it, contract negotiation is the decisive arena.

This essay focuses on smart-city initiatives—national data centres, intelligent transport, safe-city platforms, connected utilities—because these are where AI meets core public services at scale. Many examples draw from Africa-China collaborations. That is simply where a large share of deals and capital have concentrated, making the lessons practical. But the principles here are vendor-neutral and country-agnostic: they apply just as well when African governments work with partners from Europe, North America, the Gulf, or elsewhere.

What follows is a pre-negotiation playbook built for African leaders. First, we address data sovereignty: how to translate values into enforceable contract terms on localization, cross-border transfers, oversight, and rights—so sovereignty is protected in practice. Second, we turn to financial sustainability: how to structure funding for AI and smart-city projects so governments can deliver visible results today without mortgaging tomorrow—favouring transparency, lifecycle affordability, and performance-based scaling. Third, we tackle geopolitical pressure: how to keep projects resilient when global politics shift—by writing in transparency, auditability, interoperability, and the ability to renegotiate when external shocks hit.

The aim is straightforward: equip leaders to secure technology that serves public goals, strengthens institutions, and builds long-term capacity. The tools are practical—clear clauses, credible enforcement, disciplined financing, and regional coordination. With the right contracts, Africa can shape the terms of its digital transformation—and ensure smart cities are built on sovereign foundations.

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

INITIATIVE	START DATE	LEGAL NATURE	SCOPE	GOVERNING BODY
African Union Convention on Cyber Security & Personal Data Protection (Malabo Convention)	2014	Convention (requires domestic ratification & enactment)	16 AU member states have ratified the Convention	No central governing or monitoring agency
Supplementary Act on Personal Data Protection within the Economic Community of West African States (ECOWAS)	2010	Supplementary Act (binding & direct effect)	15 ECOWAS member states	ECOWAS Community Court
Southern African Development Community (SADC)’s Model Law on Data Protection	2013	Model Law (non-binding advisory instrument)		No central governing or monitoring agency
‘Support for Harmonization of the ICT Policies in Sub-Saharan Africa’ (HIPSSA)	2008	Technical assistance project		ITU and EU

NEGOTIATING DATA SOVEREIGNTY

As Africa scales AI, which can readily infer sensitive information from unsecured big data, locking in contract terms that protect citizens’ privacy and the continent’s data sovereignty is critical. Even as independent data centres expand across multiple countries (as noted in Arfan’s essay, The Rails of Inclusion: Building Africa’s Digital Foundations for AI), secure cross-border transfers and meaningful localization often lag—less because of hard infrastructure than because of soft-power gaps in law, institutions, and leverage.

Legal fragmentation of data protection regulation across the continent makes it harder to push for strict sovereignty terms with external AI service providers. Unlike the European Union (EU), the African Union (AU) does not yet have a data protection law adopted by all member states—or by all countries now moving fastest on AI. First, only 36 of 54 African countries have enacted specific data protection laws—not just data protection principles. Second, even where national laws exist, cross-border consistency is thin: continental and regional initiatives either lack collective ratification or have limited mandatory effect, as the table shows.


African leaders can advance a pan-African data protection framework by extending and strengthening existing initiatives. Draw from leading national examples (e.g. South Africa’s POPIA, Nigeria’s Data Protection Act), then—mirroring the GDPR’s status—craft a regulation with direct effect for consenting AU member states, without requiring transposition into domestic law. This boosts legal consistency and enables standard contractual clauses (SCCs) that governments can take straight to the table with third-country AI providers—just as the EU does.

Legal definition clarity can be another focus of the future pan-African data protection regulation. After all, vague definitions create wide exits for accountability.

Both the Malabo Convention (Article 1) and the GDPR (Article 4) include definitions, but Malabo is notably less precise. Firstly, the Malabo Convention omits key terms like ‘pseudonymization’, which means ‘the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information...’ in GDPR. Additionally, because the Malabo Convention is not a strictly risk-level based data protection framework, it only differentiates ‘personal data’ and ‘health data’, whereas GDPR adds the definitions of ‘genetic data’ and ‘biometric data’ on top of the previous two. Clear definitions of these terms can leave fewer grey areas for service providers and enhance monitoring, auditing, and enforcement efficiency. Secondly, African leaders must develop an original yet precise definition of ‘having adequate level of protection’, as defining adequacy is retaking agency in the data sovereignty issue. For instance, the AU can treat countries already recognized by the EU (e.g. Japan, South Korea, the United Kingdom) as presumptively adequate for specified transfers; and grant adequacy to others that sign African SCCs aligned to the new framework and meet specified safeguards. Name the categories. Set the bar. Then make every contract transfer earn it.


Independent oversight is essential to secure, lawful cross-border data flows and trustworthy AI deployment. Data Protection Authorities (DPAs) turn principles into practice: they apply data protection rules, scrutinize AI service providers, and enforce the terms leaders negotiate into their contracts. Where DPAs are credible, well-resourced, and independent, vendors take compliance seriously and citizens gain confidence that their data is protected.

Africa has already witnessed what empowered oversight can do. In 2023, Senegal’s Personal Data Protection Commission declined approval for an employer’s AI facial-recognition system in monitoring workers’ participation and then issued guidance on ethical use of workplace biometrics. That sequence—evaluate, decide, and clarify—shows how a DPA can keep AI innovation aligned with public values. Across the continent, DPAs also cooperate with peers to share information, coordinate investigations, and harmonize approaches so cross-border enforcement becomes practical rather than aspirational. This kind of cooperation lowers friction for legitimate data use while raising the cost of cutting corners. Figure 1 below, produced by policy advisory firm Tech Drive, shows interstate DPA coordination in Africa by 2023.



To meet the scale and speed of AI, DPAs need stable funding and real independence. Many authorities remain under-resourced, and some rely heavily on fines from non-compliant companies, which is not conducive to long-term planning. Multi-year public budgets, complemented by transparent supervisory fees, allow regulators to recruit, train, and retain the talent required to keep pace with fast-moving technologies. Just as important, DPAs should bring multiple perspectives to the table. Recruiting engineers, lawyers, economists, ethicists, and sector specialists—and consulting civil society, industry, labour, and academia on a regular cadence—ensures that rules reflect real-world use and anticipate unintended consequences. Finally, DPAs should be written into the life of public data and AI contracts. Designating the regulator as the monitoring authority enables regular audits, verifiable access logs for sensitive datasets, clear incident-reporting timelines, and enforceable remedies. Non-compliance must carry predictable consequences, from corrective action plans to suspension of processing to contractual penalties. At the regional level, memoranda of understanding among DPAs can enable joint investigations, shared evidence standards, and mutual assistance for cross-border transfers, giving African leaders a coherent enforcement front.

Total external debt owed by African countries



Bilateral (excl. China)
Bilateral (China)
Multilateral
Private (excl. China)
Private (China)

NEGOTIATING FINANCIAL SUSTAINABILITY

Finance will decide whether Africa’s AI transformation scales or stalls. Right now, the continent is carrying a heavy load: in 2023, external debt averaged about 24.5 percent of GDP and kept climbing; by the end of 2024, Africa’s total external debt exceeded US\$1.1 trillion. That burden doesn’t just squeeze budgets—it squeezes people. It raises the risk of debt distress, fuels poverty and inflation, erodes real income growth, and has helped spark waves of anti-government protest. This is the backdrop against which every smart-city and AI investment must be negotiated.

China remains a major creditor in this picture. As Figure 2 shows, Chinese public and private lenders accounted for roughly 12 percent of Africa’s external debt in 2023, and the share has been rising. For governments seeking to fund smart-city initiatives, that creates a real dilemma: Chinese financing can close critical infrastructure gaps quickly, yet it can also introduce repayment pressure and transparency risks. The task, therefore, is not simply to find money—it’s to secure financing on terms that power innovation today without mortgaging tomorrow.

Chinese state-backed loans are attractive for clear reasons: they come in large amounts and often with less strict and transparent political conditions than Western public finance. That combination can fund the construction of national data centres, fibre cables, and smart-city platforms at a scale private lenders rarely match. But the same features that make these loans appealing can strain long-term sustainability—confidential terms, foreign-currency repayment, and step-in or collateral provisions that narrow options if revenues underperform. On the other hand, Western financing can pose a different set of pressures. Commercial funds may be smaller, pricier, and shorter in tenor, creating refinancing necessities just as projects move from build to operate. Multilateral and bilateral financing is often cheaper but slower, with complex conditions and heavy reporting that delay disbursement and shift the focus to process over performance. Grants and pilots can launch eye-catching AI use cases without covering operations, maintenance, or data-governance costs—leaving governments to fund the lifecycle out of strained budgets. Public-private partnerships can mobilize capital but sometimes hinge on minimum-revenue guarantees and currency risks the public side ultimately bears. The result is a shared challenge across lenders: financing terms that prioritize ground-breaking announcements over durable service delivery. Negotiations for AI and smart-city projects must therefore capture the advantages—scale, speed, and affordability—while correcting structural weaknesses in transparency, currency exposure, lifecycle funding, and vendor lock-in.

Diversifying funding structures can help African leaders balance the risk of different funding sources. For example, Kenya’s Konza Technopolis, which aims to ‘develop a thriving sustainable smart city and a vibrant innovation ecosystem’, benefits from its mixed financing. Firstly, the project’s smart-city core and data centre were

funded through a US\$200.4 million concessional loan offered by China Eximbank, with 2 percent interest and generous maturity terms. Secondly, the project utilized an Italian EPC+Finance model, backed by SACE, Italy’s export credit agency, to finance its infrastructure such as roads and sewage. Thirdly, in 2024, Kenya and the Republic of Korea signed a financial agreement of US\$284.1 million towards the establishment of the Digital Media City at Konza Technopolis. And, fourthly, Kenya’s National Treasury also allocated budgets to Konza Technopolis every year: KES 5.2 billion (approximately US\$40 million) in the financial year 2024/25, and KES 3.1 billion (approximately US\$24 million) in the financial year 2025/26. This blended funding structure allowed Kenya to spread financial risk and dependence on a single creditor, thereby improving its bargaining position in smart-city contract negotiations. African leaders can therefore actively consider and search for blended financing in AI goods and services.

Additionally, it is worth noting here that Chinese AI technological goods and services can be deployed without being funded by Chinese debt. In 2018, Uganda rolled out nearly 5,500 Huawei cameras for nationwide public surveillance to enhance national security. Unusually, the project was not financed by China Eximbank but by a US\$104 million loan from Standard Chartered Bank, a British multinational bank, alongside US\$22 million from the Ugandan government budget.

Starting with smaller, repayable projects is another way to mitigate funding risk. China itself has shifted towards ‘small and beautiful’ projects under its Belt and Road Initiative (BRI)—smaller-scale, greener, and more cautious loans with clearer repayment prospects. In the Forum on China-Africa Cooperation Beijing Action Plan (2025-2027), Article 5.2.6 mentioned that China would develop a Special Loan for African small and medium-sized enterprises (SMEs) to ‘promote employment and improvement of people’s liveli-

hood through accessible and affordable financing’. It also mentions diversifying funding structures to enable different types of infrastructure projects. Thus, African leaders can encourage local businesses to start with repayable pilot projects that enable governments to test their repayment capacity, build public trust, and avoid the fiscal and reputational damage of high-profile failures.

Although not highly-AI related, the City of Cape Town’s MyCiTi bus rapid transit (BRT) system in South Africa is a real-world illustration of a phased, incremental funding strategy. Due to the budgetary constraints of local governments, it was difficult to secure the entire funding upfront. The senior leaders of the project thus broke this mega-project down into several manageable phases, starting from high-demand corridors with projects N2 Express Service and Phase 1A. These initial pilots achieved significant success, which offered the city access to funding from the national Public Transport Network Grant (PTNG) and a US\$300 million loan from the World Bank’s International Bank for Reconstruction and Development (IBRD) in 2021. This case provides a reference for African leaders, who can start AI smart-city collaborations in small, high-demand pilot areas and secure more funding after the initial success.

Technology transfer should also be prioritized in contract negotiation and drafting with external AI goods and service providers. Without knowledge transfer, smart-city projects risk leaving African countries technologically dependent on external contractors, especially when the funding source becomes unpredictable. For instance, many Education Ministries of African countries have signed agreements with Huawei to establish local Information Technology academies for technical training. The example of Senegal’s National Data Centre, built by Chinese company Huawei with Eximbank financing, also shows that local engineering training is possible if governments insist upon it. Negotiation for knowledge sharing can be based on ‘Article 5.3 Science and Technology Cooperation and Knowledge Sharing’ of the Forum on China-Africa Cooperation Beijing Action Plan (2025-2027). Actually, many African countries already have recommended Technology Transfer Regulations while dealing with foreign companies with technological advantages. Lead-

ers can make these regulations more prevalent in smart-city contracts, requiring deeper technology and knowledge transfer in exchange for access to the African market. In doing so, governments transform external finance from a dependency risk into a capacity-building opportunity.

NEGOTIATING UNDER GEOPOLITICAL PRESSURE

Geopolitically, collaboration with African countries on smart-city infrastructure has become essential for the diplomacy of superpowers. Beginning with China’s BRI, and particularly the Digital Silk Road (DSR), Beijing has invested heavily in building fibre-optic cables across more than 20 African countries to create secure communication channels independent of the US. China also has the intention to translate its economic footprint in Africa into normative influence: In September 2020, Beijing launched the Global Initiative on Data Security, which explicitly positioned China as an emerging rule-setter for international digital governance, justified on the grounds of its significant market share in AI and related technologies. Critics in the West, however, argue that such technology transfers and economic exchanges by China have geopolitical consequences. They raise concerns that safe city projects risk expanding authoritarian governance models through surveillance technology, while also entrenching Chinese monopolies over Africa’s natural resources. In this sense, African governments negotiate smart-city contracts not only on technical and financial terms but also within a larger geopolitical struggle.

African governments are not merely passive recipients of smart-city technologies but strategic actors negotiating amid great power rivalry. As such, African leaders must adopt proactive measures to safeguard sovereignty and resilience. The following recommendations consider pathways to manage geopolitical risks while maximizing developmental benefits. Many of these paths are already implemented by African countries.

Collective and regional bargaining through the AU or other regional blocs can increase Africa’s negotiating leverage and reduce fragmentation amid geopolitical turbulence. While Nilofar’s essay, *Data as Destiny: Unlocking Efficiency and Trust in Public Service*, highlights the standardization of data governance among government departments within one country, the AU’s Data

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

Policy Framework (2022) calls for facilitating ‘collaboration between the various entities dealing with data across the continent,’ supported by the AU and sister Pan-African organizations. The report strongly urges African countries to negotiate a common position among themselves for data-related international negotiations that prioritizes equal opportunity and capacity-building for developing countries. Similarly, the AU’s Digital Transformation Strategy for Africa (2020-2030) explicitly calls for harmonization of ICT frameworks and coordinated regional approaches to external partnerships. If such coordination were applied to smart-city and AI contracts, African governments could collectively resist confidentiality clauses, negotiate stronger audit rights, and ensure compatibility across borders. Acting as a bloc also makes it more difficult for external actors to play countries off against one another.

Transparency, auditability, and the potential to renegotiate should be built into every AI and smart-city contract. Contracts should include provisions that allow, to an extent, suspension, renegotiation, or substitution if geopolitical conditions shift. In practice, some external lending agreements are rigid—prohibiting renegotiation without lender approval and imposing strict confidentiality obligations. This lack of flexibility creates vulnerability: African governments can be locked into terms that become impractical if international conditions change. The risk is not hypothetical. When Huawei faced sanctions and supply-chain disruptions in 2019-2020, several African countries suddenly had critical systems dependent on service or equipment providers who were cut off from critical supply chains. Had contracts included renegotiation triggers tied to sanctions, export controls, or force majeure events, governments would have been better positioned to adapt.

Additionally, opaque financing compounds geopolitical risks. Broad confidentiality obligations that restrict disclosure—even to parliaments or independent auditors—undermine accountability and impede coordinated responses across borders. To mitigate this, African governments can insist on transparency provisions: the right of parliamentary committees and national audit offices to review loan terms; audit rights tied to project milestones; publication of non-price terms and redacted versions where appropriate; and independent verification of compliance. Transparency not only strengthens accountability but also creates space for collective bargaining, since other governments can learn from disclosed terms and align their negotiating positions. In short, write flexibility into the deal, protect interoperability from day one, and make transparency non-negotiable.

HOW AI CAN ACCELERATE INTRA-AFRICAN TRADE

OMOTAYO DADA

INTRODUCTION

What if the next great leap in African economic integration comes not from new trade agreements, but from how we apply technology to solve age-old barriers? While intracontinental trade across the world keeps increasing, formal intra-African trade remains at just 17 percent of the total continental commerce. This figure pales against Asia’s 59 percent and North America’s 51 percent. The question is not whether Africa will embrace Artificial Intelligence (AI) in trade, but whether we will shape its deployment to serve our unique contexts and aspirations.

The story of M-Pesa offers a powerful lesson in contextual innovation. When Kenya reimagined mobile banking in 2007, it didn’t simply adopt Silicon Valley’s PayPal model. Instead, it created a solution rooted in Kenya’s realities—leveraging SMS technology, working without internet connectivity, and building on networks of local agents. Within a decade, M-Pesa was processing nearly half of Kenya’s GDP. This wasn’t just technological success; it was a case of solving African problems with African ingenuity.

Today, as AI transforms global commerce, we face a moment of choice. Will we wait for AI solutions designed elsewhere to be imposed upon our markets, or will we proactively design and deploy AI solutions to address our most pressing trade challenges—from customs delays that make shipping between Lagos and Accra more expensive than shipping from Shanghai to Lagos, to language barriers that fragment our markets with over 2,000 languages?

The African Union’s Agenda 2063 envisions a continent where ‘African people have a high standard of living, quality of life, sound health and well-being.’ The African Continental Free Trade Area (AfCFTA) represents a strong commitment to economic integration, projecting growth in intra-African trade from 15 percent to 26 percent by 2045. However, achieving this goal requires

more than policy frameworks; it demands leveraging available tools, including AI, to break down the practical barriers that have limited Africa’s economic potential for decades.

It is possible for AI-powered customs systems to reduce clearance times from weeks to days across African countries. AI can help smallholder farms, which represent 80 percent of farms in sub-Saharan Africa, optimize yields and access regional markets. Digital translation tools can enable a Senegalese service provider to seamlessly work with Ethiopian clients. These are not distant dreams but achievable realities with strategic implementation.

We must approach this transformation with both ambition and wisdom. The same technology that can democratize market access can also deepen existing inequalities if deployed without careful consideration of diverse developmental realities. Countries with stronger digital infrastructure may surge ahead, while others risk further marginalization. The challenge for African leadership is ensuring that AI integration becomes a tool for inclusive growth rather than a driver of greater disparity.

This essay presents a strategic framework to harness AI’s potential in regional trade while maintaining agency over our economic destiny. It examines three high-impact areas where AI can deliver immediate value: customs process optimization, increasing agricultural productivity, and enhancement of digitally delivered services. Importantly, practical implementation principles are suggested to efficiently leverage AI for intra-African trade growth.

DESIGNING AFRICA’S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

THE AfCFTA IMPERATIVE:

WHY AI IS ESSENTIAL FOR AFRICAN TRADE SUCCESS

The AfCFTA is the world’s largest free trade area by participating countries with about 1.3 billion people and US\$3.4 trillion in combined GDP. The African Union’s projection for intra-African trade to reach 26 percent by 2045 means Africa will still lag behind Asia’s 59 percent. Are we being pragmatic or are we just content with incremental progress when transformational tools are within our grasp?

The AfCFTA aims to create a single continental market for goods and services, but it faces a paradox: reducing tariffs does not matter if trade costs make trading within the continent unattractive. Research across 22 African countries shows that reducing trade costs generates higher trade volumes than equivalent tariff cuts. You can remove every tariff, but if logistics make intra-African trade twice as expensive as intercontinental trade, the AfCFTA’s ambition would fall short.

AI can be leveraged to address key operational barriers that constrain trade. AI-powered automation is slashing trade costs globally, while generative AI (in consulting) amplifies worker performance by 40 percent, particularly for low-skilled workers. For a continent losing talent to other regions, Africa cannot afford to ignore tools that could multiply its human capital’s productivity. Notwithstanding the slow increase in the percentage of formal intra-African trade in recent years, trade volumes can still significantly increase. The AfCFTA is not just about goods; it includes services trade. Services trade, where AI’s impact is most dramatic, remains Africa’s most underexplored opportunity. Currently, Asian and European companies capture the bulk of global services trade using AI-powered platforms. The opportunity still exists for African businesses to tap into this growth.

CURRENT STATE OF INTRA-AFRICAN TRADE

In 2023, intra-African trade stood at 14.7 percent, declining marginally to 14.4 percent in 2024. In 2025, formal intra-African trade rose to 17 percent. While the share of formal trade remains lower than Asia’s 59 percent share of intracontinental trade, African countries remain optimistic that effective implementation of the AfCFTA Agreement will drive a meaningful expansion in intracontinental trade. However, for most countries, implementation of the Agreement is still at an early stage, leaving substantial room for growth as it deepens.

High trade costs have been identified as one of the primary reasons for the low share of trade between African countries. It is reported that shipping a standard 20-foot container from Shanghai (China) to Lagos (Nigeria), a distance of 20,328 km by sea, costs approximately US\$2,500–3,000, with a turnaround time of 30 days. In contrast, shipping the same container from Lagos to Accra (Ghana), a distance of less than 500 km, can cost US\$4,000–5,000 and take 7 to 10 days or more, depending on border clearance. These high-cost structures explain why non-African countries dominate Africa’s import markets. Europe and Asia remain the top import regions for African countries.

Poor infrastructure and complex trading processes have also increased trade costs. The absence of direct flights between countries, or railways to transport goods between countries, makes shipments more expensive. Goods are delayed at congested borders and ports with clearance times averaging 3 weeks in many sub-Saharan countries. The longer the shipment takes, the larger the trade costs. Thus, reducing trade costs can boost exports significantly.

Trade volumes across Africa’s regions also differ. In 2024, Southern Africa led intra-African goods trade, while North and Central Africa recorded the lowest intra-African trade shares. Southern Africa has the most substantial unrealised intra-African trade potential, estimated at US\$18.9 billion, followed by North Africa and East Africa at US\$8.7 billion and US\$7.9 billion, respectively. West Africa’s unrealised potential stands at approximately US\$4 billion, while Central Africa’s is around US\$2.3 billion. This highlights both the opportunity and the uneven distribution of trade capacity across the continent.

Africa's path towards AI-powered trade solutions faces additional complexities. While countries like South Africa, Egypt, Rwanda, Mauritius, and Nigeria have emerged as frontrunners in AI governance, with Kenya leading as the 'Silicon Savannah' in innovation and policy development, the continent's AI readiness remains highly uneven. The digital divide between countries poses a particularly acute challenge, threatening to transform AI from a tool of integration into a driver of further fragmentation. Countries with robust internet infrastructure, reliable electricity grids, and advanced telecommunications networks can rapidly deploy AI solutions for customs automation, digital services platforms, and agricultural optimization. Meanwhile, nations struggling with intermittent connectivity, limited bandwidth, and inadequate digital infrastructure find themselves locked out of AI-enabled trade opportunities entirely. Countries lacking AI infrastructure risk becoming increasingly marginalized in global markets, creating new forms of trade inequality that could deepen existing regional disparities.

Critically, the rapid pace of AI development outstrips traditional trade governance mechanisms. This mismatch between technological capability and regulatory frameworks creates uncertainty and potential conflicts that could undermine AI integration benefits. This is why the implementation of the Africa AI Strategy 2024 is important. The strategy establishes a multi-tiered governance approach and calls for regional coordination to help Africa build capacity collectively.

However, with the implementation timeline running from 2025-2030, significant acceleration of efforts is required to bridge the gap between the strategy's ambitious vision and the reality on the ground. Each African country must chart its own AI course while contributing to a collective vision. Harmonization efforts are crucial for avoiding policy fragmentation and duplication and for supporting the African Continental Free Trade Area, as well as providing the continent with a unified and influential global voice.

USE CASES FOR AI INTEGRATION

Several segments of intercontinental trade offer opportunities for AI integration, but the focus in this essay is on three areas: (i) customs process optimization, (ii) agricultural productivity, and (iii) provision of digitally delivered services. The basis for this selection is the African Union’s Agenda 2063 as well as the strategic importance of these areas. They target immediate trade constraints while laying the foundation for sustainable economic growth. Africans still rely on roads to transport around 80 percent of goods, and border inefficiencies continue to drive delays, raise costs, and reduce trade volumes. Optimizing customs processes through AI therefore represents a critical intervention. In agriculture, which employs about 46 percent of Africa’s workforce, AI can enhance food production, improve food security, and generate tradeable surpluses. Finally, while services trade contributes significantly to Africa’s GDP, the continent’s share of global services trade is less than three percent. AI can be leveraged to increase this, allowing Africa to tap into the sector’s expanding potential.

These three areas create mutually reinforcing benefits. They support the aspirations of the African Union’s Agenda 2063 for economic transformation, technological advancement, and continental integration while simultaneously addressing the immediate constraints that limit Africa’s participation in global trade.

CUSTOMS PROCESS OPTIMIZATION

Goods trade dominates African commerce, making customs and port authorities the gatekeepers of regional integration. Currently, customs clearance averages 12 days in sub-Saharan Africa—nearly double the seven days in Latin America and more than double Central Asia's five and a half days. Every delayed shipment represents not just higher costs, but lost competitiveness and frustrated traders who increasingly look elsewhere for markets.

AI can transform this chokepoint into a competitive advantage through four critical interventions:

1.
Automation and Streamlining:
Intelligent document processing and automated compliance checks can eliminate manual bottlenecks that currently trap shipments for weeks. While this creates short-term workforce displacement, strategic retraining programmes can redirect customs officers towards high-value analytical tasks that AI cannot perform.
2.
Revenue Collection and Self-Service:
Mobile payment platforms like Botswana's BURS Mtax and Zambia's TaxOnApp demonstrate how AI-enabled systems can streamline duty collection while improving transparency and reducing opportunities for corruption.
3.
Real-Time Risk Assessment :
Nigeria's Integrated Customs Information System (NICIS) showcases AI's potential to process massive data volumes instantly, enabling officers to focus inspections on genuine risks rather than conducting time-consuming blanket checks.
4.
Customer Care Optimization :
AI-powered chatbots can handle routine inquiries and guide traders through complex procedures, freeing human officers for exceptions that require judgment and experience.

The evidence of the impact that can be achieved is compelling. Research across 22 African countries demonstrates that trade cost reductions generate greater increases in trade volumes than equivalent tariff cuts. High trade costs may indeed be a worse enemy than tariffs themselves.

AGRICULTURAL PRODUCTIVITY

Mainstreaming AI in a priority sector like agriculture aligns with the African Union’s agenda for a prosperous Africa based on inclusive growth and sustainable development. Agriculture employs about 46 percent of Africa's workforce, making technological advancement in this sector essential for both food security and export competitiveness. 33 million smallholder farms represent 80 percent of sub-Saharan Africa's agricultural operations.

The global AI agriculture market is projected to grow at 23 percent annually from US\$1.7 billion to US\$4.7 billion between 2023-2028. This is an opportunity that Africa cannot afford to miss out on. Early implementations already demonstrate transformative potential:

1.
Hello Tractor's platform has digitized 3.5 million acres since 2014, increasing food production by 5 million metric tons while creating over 6,000 jobs. The platform uses machine learning to optimize tractor deployment, monitor usage patterns, and enable SMS-based coordination in areas with limited internet access.
2.
Kenya's Agricultural Observatory Platform delivers real-time data to 1.1 million farmers, providing accurate weather forecasts and agricultural insights that optimize planting and harvesting decisions. The World Bank's expansion of this technology through the Food System Resilience Program now benefits 6 million farmers across West Africa.

3. Cameroon's AI-powered disease identification app enables farmers to photograph affected crops and receive instant diagnoses with treatment recommendations. An important feature of this is that it works offline to serve areas without reliable internet connectivity.

It is clear that AI can be leveraged to improve productivity in agriculture for smallholder farmers in Africa. Designing and training AI models requires large and accurate datasets; without them, fragmented African datasets risk producing systems that perform poorly in our local contexts.

DIGITALLY DELIVERED SERVICES

Global services trade outpaced goods trade growth in the third quarter of 2024, offering African countries the opportunity to participate in high-value commerce without traditional infrastructure requirements. For countries constrained by geography or resource limitations, digital services represent a pathway to economic relevance that bypasses many conventional barriers. Four strategic applications can position Africa as a services trade competitor:

1.
AI-Powered Market Matching:
Intelligent platforms can connect African service providers with continental and global buyers, reducing search costs and improving match quality. OECD research demonstrates that AI-enabled matching delivers both efficiency gains and cost savings while accessing larger talent pools.
2.
Digital Service Innovation:
AI enables entirely new categories of tradeable services, from AI-powered analytics to intelligent software solutions and virtual professional services. These require internet connectivity rather than physical infrastructure, making them accessible to countries across Africa's developmental spectrum.

3. **Productivity Amplification:**
BCG's 2023 research found that generative AI amplifies consultant performance by up to 40 percent. More significantly, a customer service study revealed 14 percent average productivity gains, rising to 34 percent for novice and low-skilled workers. This suggests AI could democratize access to high-productivity work patterns, particularly benefiting countries with emerging skill bases.

4. **Elimination of Language Barriers:**
Translation technology can unlock continental market access for service providers. While Africa's 2,000 languages across 54 countries present daunting complexity, a phased approach of designing translation applications for major trade languages like Swahili, Hausa and Arabic can deliver immediate benefits while building towards comprehensive coverage. Research has found that reducing language barriers has a positive impact on trade.

Basic digital skills are sufficient to use AI-enabled services, but African countries seeking to develop or adapt AI models will require a workforce with advanced competencies in information technology and machine learning. These capabilities not only support domestic innovation but also attract foreign investment and facilitate integration into the rapidly growing digital services trade.

These three AI application channels are not independent; they are mutually reinforcing components of a comprehensive trade ecosystem transformation. Customs optimization reduces trade costs. Agricultural productivity creates tradeable surpluses. Digital services expand participation opportunities for countries regardless of natural resource endowments.

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

WHY AI-ENHANCED TRADE TRANSCENDS
ECONOMIC BENEFITS

AI-powered trade integration delivers returns that extend beyond GDP growth; it reshapes Africa's political landscape and strategic sovereignty.

1.
Political Leverage and Regional Stability: Economic interdependence through AI-enhanced trade networks creates powerful incentives for peaceful conflict resolution. When countries depend on each other for critical goods and services flowing through automated systems, diplomatic stability becomes an economic necessity, not just a political preference.
2.
Food System Resilience:
AI-optimized agricultural networks enable real-time coordination between surplus and deficit regions, transforming food crises into manageable distribution challenges. Smart trade systems can predict shortages before they become famines, and supplies can be redirected to prevent humanitarian disasters.
3.
Indigenous Innovation Ecosystems:
Enhanced trade relationships become conduits for sharing technological breakthroughs, agricultural innovations, and business models between African countries. This reduces our dependence on external knowledge sources while building indigenous capabilities that serve African contexts.
4.
Strategic Autonomy:
Robust intra-African trade reduces our vulnerability to external economic manipulation and political pressure. A continent that trades primarily with itself cannot be easily coerced through sanctions, trade restrictions, or commodity price manipulation by external powers.

CRITICAL IMPLEMENTATION PRINCIPLES

Half-hearted AI adoption will create the worst of both worlds: job displacement without productivity gains, digital divides that fragment our markets further, and technological dependence that locks us into permanent subordination.

Moving from AI experimentation to meaningful impact requires clear implementation principles, like those captured below, that are grounded in Africa's realities.

1.
Leverage existing successes:
Rather than creating new solutions, prioritize scaling proven African AI applications like Hello Tractor, NICIS, and the Agricultural Observatory Platform.
2.
Sequence strategically:
Ensure foundational digital infrastructure exists before deploying sophisticated AI applications, avoiding the trap of advanced technology on inadequate foundations. As outlined in Arfan's essay The Rails of Inclusion: Building Africa's Digital Foundations for AI, there is a need to 'build the rails' for inclusive and effective solutions.
3.
Address digital divides proactively:
Provide technical and financial support to help less-developed countries participate in AI-enabled trade networks from the beginning, preventing further marginalization.
4.
Prioritize continental coordination:
Use AfCFTA protocols, particularly the Digital Trade Protocol, as frameworks for harmonizing AI implementation across borders rather than pursuing fragmented national approaches.

CONCLUSION

It is impossible for countries to all move at the same implementation speed; some countries on the continent can lead and become major ‘AI trade hubs’ and then pull others along. Africa's AI-enabled trade integration journey can draw valuable lessons from other developing regions navigating similar transitions. Southeast Asia's experience with the ASEAN Single Window, which uses AI-powered customs systems across ten diverse economies, demonstrates that coordinated digital integration is achievable even amid significant developmental disparities. Thailand reduced customs clearance times from seven days to under 24 hours through phased AI adoption, while Singapore's role as a digital hub created spillover benefits for neighbouring countries through technology transfer and skills development.

Countries that effectively deploy AI across customs systems, agricultural systems, and digital services trade today will capture exponential advantages tomorrow. Those that hesitate will watch trade flows bypass their borders, agricultural productivity stagnate, and service opportunities migrate to better markets. The mathematics are brutal; in a world where AI amplifies worker productivity by 40 percent and slashes trade costs globally, standing still is detrimental. Investment in digital infrastructure and workforce skills is increasingly recognized as a key driver of comparative advantage in the emerging AI-driven global economy.

Beyond economic benefits, it is about African agency. Every month we delay implementing AI trade solutions is another month of shipping goods from Shanghai to Lagos more cheaply than from Lagos to Accra. Every quarter we spend debating instead of deploying is market share permanently lost to continents that act with urgency.

The African Union's vision of 26 percent intra-African trade by 2045 is achievable if we embrace AI as aggressively as we have embraced mobile technology. The path forward demands neither haste nor hesitation, but purposeful action.

CONCLUSION

AARON MANIAM

Countries in Africa face a moment of both extraordinary promise and deep structural risk in how they manage Generative AI. The continent has the opportunity to leapfrog existing technologies, using AI to improve public service delivery, expand access to markets, accelerate scientific discovery, and strengthen cultural production. At the same time, as in other parts of the world, AI threatens to entrench exclusion, extract value without building capacity, and automate existing inequities at unprecedented speed.

The essays in this collection converge on a central insight: AI outcomes are determined far less by algorithms than by institutions, incentives, and infrastructure. Identity systems, data governance, contracts, human capability, and social legitimacy – all intrinsically human issues – will matter more than model size or compute power. Generative AI will not transform Africa by default. It will do so only if transformation is intentionally designed, collectively governed, and faithfully implemented.

This essay builds on the many insights in this collection, extending them into five practical steps that governments, businesses, and NGOs can take together to ensure that Generative AI becomes a force for inclusive development rather than digital exclusion. In each section, concrete next steps are identified for different actors in multi-stakeholder systems.

**STEP 1:
ANCHOR AI TRANSFORMATION IN VISION, STAMINA,
AND IMPLEMENTATION**

Transformation begins with vision—but it succeeds or fails through execution. Without sustained political commitment and institutional discipline, AI adoption will simply mirror existing bureaucratic cultures, power structures, and inequalities. As Conway’s Law reminds us, systems reflect the organizations that build them. If ministries are fragmented, AI systems will be fragmented. If incentives reward opacity, AI will

scale opacity. If elites dominate decision-making, AI will accelerate elite capture.

Across Africa, there is no shortage of AI strategies, task forces, and pilot projects. These must be complemented by fidelity of implementation: the ability to follow through over multiple budget cycles, leadership changes, and political shocks. Arfan Ramadhan’s analysis of digital public infrastructure makes clear that inclusive outcomes depend on unglamorous but critical long-term investments in identity, payments, and data exchange. Omotayo Dada’s exploration of the impact of AI on intra-African trade similarly emphasizes the criticality of effective implementation to realize the benefits of customs process optimization, agricultural productivity, and digitally delivered services. The foundations of such technological benefits take years, not months, to mature, and will require intrinsically long-term outlooks.

CONCRETE NEXT STEPS

Governments should treat AI and digital foundations as core state capacity, not innovation side-projects. This means anchoring them in central ministries with budget authority, audit requirements, and long-term mandates.

Businesses should resist ‘AI-washing’ and align their deployments with institutional realities. This will include designing for interoperability, resilience, and public accountability, building on but also moving beyond short-term demonstration projects.

NGOs should act as implementation watchdogs, tracking whether announced AI initiatives actually deliver inclusive outcomes on the ground.

DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value

**STEP 2:
MAKE AI A MODERN UBUNTU**

Generative AI is increasingly eroding traditional sectoral boundaries. A single use case, like AI-enabled agricultural advice through an app, touches telecommunications, education, finance, data protection, media, and local governance simultaneously. The essays by Sitong Lyu and Siaka Salami show that siloed approaches are not merely inefficient; they are actively dangerous, producing brittle systems that collapse under the weight of real-world complexity.

Africa’s response must therefore be explicitly multi-sectoral. Governments, businesses, NGOs, academics, journalists, artists, and community leaders must co-govern AI systems that shape public life. This is not a romantic appeal to consultation or superficial collaboration, but a practical necessity for legitimacy and resilience.

Crucially, citizens must be engaged not only as ‘users’ but as rights-bearing participants in decisions about trade-offs. Marlon-Ralph Nyakabau’s human-centric framework warns against equating convenience with dignity. Generative AI raises real risks: cognitive offloading that erodes skills, misinformation at scale, automated exclusion, and surveillance creep. These trade-offs cannot be resolved technocratically but require deep democratic deliberation.

This is the oft-cited African philosophy of Ubuntu in its modern form: ‘I am because we are’, extended to digital systems. Done well, Gen AI will bring new dimensions and richness to this idea.

CONCRETE NEXT STEPS

Establish permanent, multi-stakeholder AI councils at national and city levels, with real influence over procurement, standards, evaluation, policymaking, and input into legislative discussions.

Require meaningful public consultation for high-risk AI deployments, including clear explanations of benefits, risks, and safeguards.

Support media literacy, constructive journalism, and academic research as integral components of AI governance – not external critics.

**STEP 3:
TREAT PEOPLE AND TALENT
AS THE CORE CONSTRAINT**

Across all this collection’s essays, one conclusion is unavoidable: talent is the rate-limiting factor. Data infrastructure can be procured; models can be licensed; cloud and compute capacity can be rented (admittedly with tradeoffs to sovereignty – but some of this can be managed, especially in domains where open platforms carry acceptable risk). But without skilled people to design, govern, audit, and adapt AI systems, Africa will remain dependent on external vendors and external priorities.

Nilofar Saiyed highlights how weak data ecosystems stem not from technology gaps but from institutional and human capacity gaps: analysts cleaning poor data, civil servants lacking incentives to share information, leaders without real-time insight into implementation. Sitong extends this to contract negotiation: without skilled lawyers, engineers, and economists, African governments negotiate from weakness, locking in vendor dependence. African countries therefore need a dual talent strategy: fair access to global talent flows, especially to meet short-term needs; and deep investment in home-grown capability, to prepare for the long-term transformative effects of AI.

CONCRETE NEXT STEPS

Governments should create competitive career tracks for data, AI, and digital policy professionals within the civil service, including flexible pay and secondments. This could even extend to a ‘Digital Public Service Corps,’ recruiting young engineers, designers, and policy graduates into three-year rotations across ministries - similar to the Peace Corps and ‘Teach for America’ programmes in the United States. Fellows receive advanced AI training, but also work on procurement, data governance, and citizen engagement - creating a pipeline of leaders who understand both technology and the state.

Businesses should invest in local R&D, apprenticeships, and open-source ecosystems rather than relying solely on imported solutions.

NGOs and universities should focus on interdisciplinary training - combining AI with law, ethics, public policy, and social sciences.

**STEP 4:
COMBINE GENERATIVE AI WITH DIGITAL
PUBLIC INFRASTRUCTURE**

Generative AI delivers its greatest public value when layered onto Digital Public Infrastructure (DPI): identity systems, interoperable data exchange, and payment platforms. Arfan’s ‘rails of inclusion’ argument is instructive here – AI cannot target, personalize, or scale benefits for people who are invisible to the system. DPI turns AI from an abstract capability into a practical instrument of inclusion. Identity enables recognition; data exchange enables coordination; payments enable delivery to meet human needs. When these layers are designed with consent, interoperability, and resilience, AI can move from pilots to population-scale impact.

However, combining AI and DPI also multiplies risk. Errors propagate faster. Exclusion scales more efficiently. That is why design choices – offline functionality, minimal data collection, audit logs, and grievance redress – are not technical details but decisions with significant moral dimensions.

CONCRETE NEXT STEPS

Governments should explicitly align AI strategies with DPI roadmaps, ensuring shared standards and governance.

Businesses should design AI products that plug into public infrastructure via open APIs rather than proprietary lock-ins.

NGOs should focus on last-mile inclusion: helping marginalized groups enrol, understand, and contest AI-mediated decisions.

**DESIGNING AFRICA'S
INTELLIGENCE FUTURE:
Essays on AI, Governance,
and Public Value**

**STEP 5:
BUILD STRONG DATA POLICY, REGULATION,
AND RECOGNITION FRAMEWORKS**

Data is the fuel of Generative AI, and hence both a source and terrain of power. Without clear rules, AI systems drift towards surveillance, extraction, and abuse. The essays collectively argue that good data governance is not a brake on innovation; it is a precondition for trust and scale. Sitong’s analysis of contract negotiation shows how sovereignty is exercised through enforceable clauses on data ownership, localization, audit rights, and exit options. Nilofar demonstrates that without integrated, high-quality data, policy becomes guesswork. Siaka warns that design without public value orientation erodes legitimacy, even when systems ‘work’. Omotayo’s essay argues that robust systems interoperability between countries is the bedrock for mutual trade and exchange of goods and capital, and usefully reminds us that such exchange can catalyse equally important flows of ideas, knowledge, and expertise.

Regulation must therefore do three things simultaneously:

1. Enable secure, interoperable data sharing.
2. Protect privacy, dignity, and due process.
3. Anticipate and mitigate systemic and existential risks from AI misuse.

CONCRETE NEXT STEPS

Enact and enforce comprehensive data protection laws, backed by independent regulators with real capacity.

Require algorithmic impact assessments, auditability, and human-review rights for high-risk AI systems.

Promote data recognition frameworks that value local languages, cultural content, and indigenous knowledge as assets - not noise.

**CONCLUSION:
DISCERNMENT, NOT FEAR; AMBITION, NOT NAIVETE**

Africa does not face a binary choice between embracing Generative AI and protecting its people. The real choice is between intentional, values-driven transformation and passive adoption on others’ terms. The five steps outlined here – vision with execution, multi-sector collaboration, talent investment, AI-enabled public infrastructure, and robust data governance – form a coherent path towards seizing AI’s upside while managing its risks.

The opportunity is immense. But so is the responsibility. As an African proverb reminds us: ‘The axe forgets; the tree remembers.’

Generative AI can be a powerful axe – shaping economies, states, and societies at speed. The task for leaders in Africa is to ensure that, in the rush to carve a path to the future, it does not forget the people, values, and institutions that must live with the consequences. Discernment, guided by Ubuntu and human-centred technology, is what will turn opportunity into shared prosperity.

Pax Technologica

Pax Technologica is a global initiative dedicated to radically improving the adoption of exponential technologies, particularly artificial intelligence, in emerging geographies of the Global South. We envision a harmonious relationship between humanity and technology, which will enable a future of peace, equity, and cultural enrichment. Recognizing that the next wave of global growth will be shaped by regions like the African continent, we focus on empowering local leaders, innovators, and communities to leapfrog outdated systems and harness technology in ways that are ethical, sustainable, and culturally aligned.

Pax Technologica was founded in response to the significant challenges posed by the rise of exponential technologies, particularly artificial intelligence. The troubling events of 2016 – where technology played a critical role in threatening the integrity of the US elections – served as an early warning of developments that may now accelerate rapidly. Yet alongside these challenges lies a landscape of opportunities. While many initiatives focus primarily on mitigating risks, Pax Technologica prioritizes identifying opportunities where the benefits of thoughtful and appropriate technology adoption outweigh the risks.

Through initiatives such as the Fellowship Program, Policy Intelligence Unit, strategic salons, and curated dialogues, Pax Technologica brings together diverse voices to co-create solutions for critical challenges in climate, healthcare, education, and food systems. Across all our activities, we strive to address the inequities of the digital divide while inspiring a global conversation about technology’s role in fostering a more inclusive and balanced world.

Pax Technologica has its roots in creativity, drawing from our backgrounds as designers, artists, urbanists, and storytellers. We believe in the transformative power of narrative and beauty to shape behavior and drive the changes we seek in the world. As such, Pax Technologica is as much about exploring the narrative and aesthetic dimensions of exponential technology and societal transformation as it is about making tangible change happen on the ground.

