

Policy Brief No. 239 – June 2026

Beyond Harmony, Beyond Rhetoric: Politicizing Ubuntu for AI Governance

Anye-Nkwenti Nyamnjoh

Key Points

- In artificial intelligence (AI) ethics, Ubuntu is increasingly invoked as a relational African ethical resource and as an African alternative to liberal individualism, signalling epistemic presence without necessarily reshaping how harms, power and accountability are handled in AI governance — a pattern that can risk ethics washing and leave existing governance failures intact.
- Moving beyond rhetoric requires a political turn in how Ubuntu is read and applied. A political reading treats AI systems as socio-technical infrastructures embedded in relations of power, and understands governance failures as products of institutional arrangements that concentrate authority, externalize harm and limit contestation.
- If Ubuntu is to matter in practice, it must have institutional consequence. This policy brief translates a politicized Ubuntu into concrete governance moves — participatory authority, reciprocity in data practices, restorative responses to harm and collective data stewardship — without requiring Ubuntu’s formal adoption as a policy framework.

Introduction

This policy brief is a translation exercise. It starts from the observation that Ubuntu has become a prominent ethical resource in academic AI ethics but remains thinly specified in policy discourse and rarely translated into governance design. This brief’s central argument is that the move from ethical vocabulary to policy language is not neutral and that unless questions of power, institutional design, contestation and remedy are brought into view, Ubuntu risks entering AI governance as rhetorical affirmation without being meaningfully translated into a basis for institutional design.

Ubuntu is part of a wider family of African relational moral philosophies, formalized in African philosophical scholarship, that ground personhood and ethical life in social relations and community (Molefe 2025). It is this more formalized discourse, often discussed alongside Afro-communitarianism, that has become a prominent reference point in AI ethics, especially in efforts to Africanize ethical reflection on emerging technologies (Gwagwa, Kazim and Hilliard 2022; van Norren 2023). Africanization is an interpretive lens rather than an official doctrine. It names a demand that AI ethics and governance be adapted to African contexts rather than imported uncritically from frameworks developed elsewhere.

About the Author

Anye-Nkweni Nyamnjoh is a senior research officer at the University of Cape Town, South Africa, where he coordinates the EthicsLab's research development program. His work sits at the intersection of African political thought, ethics and emerging technologies, and his scholarship contributes to debates on decolonization as an ethical, intellectual and political project. His research has explored how questions of knowledge, identity and belonging are articulated and contested within African political and intellectual traditions, and how these dynamics shape ethical and political claims about justice both past and present. He is particularly interested in how African thought is leveraged in contemporary debates on AI ethics.

Ubuntu and Afro-communitarian thought are ethical resources through which scholars articulate alignment in terms of relationality, interdependence and shared responsibility (Segun 2024; Mensah and Van Wynsberghe 2025). Relatedly, Ubuntu is sometimes mobilized as a decolonial idiom, situating AI ethics within post-colonial conditions of epistemic marginalization and structural inequality (Mohamed, Png and Isaac 2020; Adams 2021).

A parallel demand for endogeneity is evident in national and continental AI strategies, which emphasize systems that are contextually appropriate, locally relevant and aligned with development priorities. While these policy documents rarely use the language of Africanization explicitly, they often pursue a material version of it — framing alignment in terms of development, capacity building, competitiveness and the position of African countries within global AI value chains (African Union Commission 2024).

Yet Ubuntu's presence in policy discourse remains thin. Ethical traditions may be gestured toward as "African values," but invocations are often vague and underspecified, functioning more as a rhetorical affirmation than as a substantive guide to governance design. Empirical analysis of national and continental AI strategies similarly suggests that Ubuntu is rarely operationalized in

policy design, with many frameworks reproducing globally circulating ethical templates rather than articulating a grounded African ethical orientation (Jobin, Ienca and Vayena 2019; Yilma 2025).

This brief therefore addresses a pre-policy risk. In academic and normative AI ethics discourse, Ubuntu is increasingly invoked to signal African ethical presence but often in forms ill-equipped to engage power, political economy and institutional change. The danger is that Ubuntu may travel into governance as values talk: broad appeals to harmony, community or solidarity that displace harder questions about who benefits from AI systems, who bears their costs and how power is exercised.

What emerges across these opening distinctions is that Ubuntu is operating in three related but distinct registers: as a moral-philosophical discourse, as an ethical resource in AI ethics scholarship and as a thin or uneven reference in policy discourse. The concern of this brief lies in the movement between the latter two registers. More specifically, it highlights the risk that Ubuntu may enter governance discourse as ethical vocabulary without being meaningfully translated into a basis for institutional design.

Recent Ubuntu-in-AI-ethics work offers substantive conceptual reframing and power-attentive critique but often stops short of specifying concrete governance instruments (Reviglio and Alunge 2020; Mensah and Van Wynsberghe 2025). Building on this work, the brief takes a political reading of Ubuntu and translates it into implementable governance moves that can be adopted within existing policy and regulatory architectures, without requiring Ubuntu's formal adoption as a policy framework.

Ubuntu in AI Ethics: From Ethical Resource to Rhetorical Gesture

Ubuntu's prominence in AI ethics is tied to its role in asserting African ethical presence within a field shaped by global inequalities in knowledge production (Segun 2021; de Vries 2024). It provides resources for challenging the assumption that

ethical frameworks developed elsewhere can be applied uncritically across contexts and for reframing who speaks authoritatively about AI ethics and how ethical questions are situated within African intellectual traditions and social realities (Eke and Ogoh 2022).

In much of the literature, Ubuntu is invoked as a relational framework that foregrounds community, interdependence, shared responsibility and the moral significance of social relationships (Ewuoso 2021). It offers a vocabulary for reframing ethical questions away from isolated individuals and toward collective well-being and mutual recognition (Mhlambi and Tiribelli 2023). This move has intuitive appeal where AI systems shape access to public services, employment, credit, education and security.

In practice, however, Ubuntu is sometimes mobilized more as an ethical orientation than as a governance method. In some Ubuntu-inspired AI ethics work, this takes the form of appeals to community involvement, openness, harmony, support or shared responsibility as ethical ideals, without equivalent clarity about how such commitments should be institutionalized, enforced, assessed or contested (Ferlito et al. 2024; Odera, Nderitu and Samuel 2024). Ubuntu is thus presented as a set of values that AI systems should respect or reflect but often without clarity about whose interpretations matter, how conflicts between interests should be adjudicated or what institutional consequences follow when those values are violated. Ethical concern is expressed in broad terms, while governance questions — who decides, who benefits, who bears risk and who can contest outcomes — remain underspecified. In this form, Ubuntu can function as a rhetorical gesture that affirms African ethical presence while doing little to disrupt prevailing governance arrangements.

A central limitation of harmony-centred invocations is their tendency to treat conflict as ethical failure rather than as a structurally normal feature of political and economic life. Yet many AI governance challenges are irreducibly antagonistic. Platform operators exercise algorithmic control over workers with limited bargaining power; state actors deploy automated surveillance with few avenues for contestation; and extractive actors capture data, labour or resources while communities bear social and environmental harms (Crawford 2021). These relationships are sustained by institutional

arrangements and political economies, not by misunderstanding or lack of ethical intent.

In such contexts, a premium on harmony can obscure antagonisms rather than clarify what justice requires. Ethical discussion may emphasize inclusion or shared humanity while leaving intact the structures through which AI systems concentrate power, externalize costs and reproduce inequality. Ubuntu then becomes vulnerable to ethics washing (Bietti 2021).

This brief uses ethics washing in a structural rather than intentional sense (Biddle, Nelson and Olugbade 2024). The issue is not whether actors sincerely hold ethical commitments but whether ethical discourse constrains power, redistributes agency and produces enforceable outcomes. Ethics washing, on this account, occurs when ethical language legitimizes AI governance arrangements without altering the institutional conditions through which harm is produced, a recurring critique of AI ethics as depoliticized and weakly enforced (Hagendorff 2020).

The challenge, then, is not to abandon Ubuntu as an ethical resource, but to clarify the conditions under which it can inform AI governance substantively. This requires moving beyond value affirmation toward a political reading of Ubuntu — one that treats AI systems as instruments of power that reorder social relations and that demands corresponding forms of accountability, participation and institutional change (van Maanen 2022).

What a Political Turn in Ubuntu Requires

A political reading of Ubuntu treats AI systems as socio-technical arrangements embedded in relations of power. “Political,” here, means more than social embeddedness. It refers to the organization of power through institutions, infrastructures and decision-making procedures that distribute authority, risk, benefit and the capacity to contest outcomes. Ethical evaluation therefore turns on consequence and contestation: who gains control, who bears risk and what mechanisms exist to challenge, constrain or redistribute decision-making authority (ibid.).

A political turn also treats conflict as normal rather than exceptional. Many AI governance challenges are sustained by institutional arrangements and political economies rather than by the absence of ethical intent (Birhane 2021). Appeals to harmony are therefore insufficient on their own. A politicized Ubuntu directs attention to antagonisms and demands governance responses capable of addressing them through enforceable rules, institutional oversight and meaningful participation by affected groups.

Finally, a political turn redefines what counts as ethical success. It is not satisfied by compliance with principles or mitigation of isolated risks but asks whether governance contributes to justice, relational repair and the flourishing of those most burdened by inequality (ibid.). Ubuntu need not be adopted as a formal policy framework for this to matter; what is required are mechanisms that deliver these outcomes in practice.

Governance Options and Trade-Offs

Against this backdrop, it is worth considering how existing governance approaches perform when assessed against the demands of a political, power-attentive reading of Ubuntu. This section evaluates three governance logics currently shaping AI policy, examining what each enables — and where each falls short — when judged against the standards set out above. The three logics below are not mutually exclusive; the decisive question is which logic does the governing work.

Option 1: Values- and Principles-Led Governance

A first and still dominant approach relies on high-level ethical principles such as fairness, transparency, accountability and human rights (Jobin, Ienca and Vayena 2019; Hagendorff 2020). This logic is reflected in widely cited international instruments, including the Organisation for Economic Co-operation and Development's Recommendation on Artificial Intelligence and the United Nations Educational, Scientific and Cultural Organization's Recommendation on the Ethics of Artificial

Intelligence, which provide normative guidance without imposing binding legal obligations (Organisation for Economic Co-operation and Development 2025; United Nations Educational, Scientific and Cultural Organization 2021).

Principles-led governance has advantages. It establishes a shared vocabulary that supports coordination, public justification and cross-sector dialogue. It also lowers barriers to participation, as governments and organizations can signal ethical commitments quickly without extensive regulatory infrastructure.

However, principles alone rarely specify how ethical commitments translate into institutional practice. They tend to leave unresolved questions of authority, enforcement, contestation and remedy — precisely what matters when AI systems generate harm under conditions of structural inequality. As a result, principles-led governance is vulnerable to ethics washing, whereby ethical language can confer legitimacy without constraining decision making or redistributing power (Hao 2019; Rességuier and Rodrigues 2020).

Viewed through a political Ubuntu lens, this approach treats relationality as a moral orientation rather than as a governance demand. The trade-off is speed and symbolic legitimacy versus enforceability. In practice, this approach tends to benefit actors who prefer flexibility — often firms and state agencies — while offering limited leverage to workers, communities and citizens affected by AI-driven decisions.

Option 2: Procedural and Risk-Management Governance

A second approach operationalizes ethics through procedural requirements, including risk assessments, impact assessments, audits, documentation and transparency obligations across the AI life cycle. This logic underpins the National Institute of Standards and Technology's AI Risk Management Framework and a growing set of governance tools that require structured assessment of automated decision-making systems (National Institute of Standards and Technology 2023).

Procedural governance offers important gains. By embedding ethics into repeatable processes, it creates artifacts that can be inspected, audited and contested. The African Union's Continental Artificial

Intelligence Strategy is positioned as a broad governance instrument that calls for institutional mechanisms and robust ethical regulation (African Union Commission 2024). In Kenya, the litigation of the Huduma Namba digital identification cards has been cited as an illustration of how data-protection impact assessment requirements can operate as enforceable checks that must be satisfied before rollout (Muya 2021). Rwanda's data protection framework likewise operationalizes accountability through record-keeping obligations and protections relating to decisions based solely on automated processing.¹

Yet this approach has a characteristic limitation. Risk-management frameworks can frame harm as a technical or managerial problem, displacing questions about institutional design, political economy and the distribution of power. In such settings, governance prioritizes optimization, monitoring and documentation rather than substantive justice. Compliance risks becoming procedural, with assessments conducted internally and transparency offered without corresponding rights to contest decisions, seek remedy or refuse participation. The result is often an expansion of administrative visibility and control that leaves underlying power relations largely intact (Yeung 2018).

From a politicized Ubuntu perspective, procedural governance is therefore ambivalent. It can create concrete hooks for identifying relational harms but remains politically thin unless tied to enforceable constraints and avenues for contestation. The trade-off is administrative feasibility and scalability versus depth of power redistribution.

Option 3: Structural and Political-Economy Governance

A third governance approach, centred on structural political economy, starts from the recognition that many AI harms are produced by institutional arrangements and business models built around extraction, surveillance and control. Structural governance targets the conditions under which AI systems are developed and deployed, focusing on labour relations, data value chains, surveillance infrastructures, environmental externalities and

market concentration (Couldry and Mejias 2019; Crawford 2021).

This logic is difficult to locate in existing binding regulatory initiatives, which tend to engage structural concerns only indirectly or at the level of downstream effects. The European Union's Artificial Intelligence Act, for example, introduces a risk-based legal framework that prohibits certain practices and imposes enforceable obligations on high-risk systems, including requirements related to oversight, documentation and fundamental rights protection.² While these measures acknowledge that AI harms are often institutionally produced, they largely regulate the conditions of use rather than the political-economic arrangements — such as data extraction, platform business models or market concentration — that shape how AI systems are developed and scaled.

Structural governance is demanding. It requires regulatory capacity, political will and a willingness to confront resistance from powerful corporate and state actors. It can slow deployment and increase compliance costs. But these costs are not incidental as they often reflect the difficulty of governing antagonistic relations rather than merely managing reputational risk.

This approach aligns most closely with a political reading of Ubuntu, treating AI governance as the design of just relations under conditions of unequal power rather than as the articulation of shared values. It is therefore best positioned to translate Ubuntu's ethical demands — relational repair, accountability and the flourishing of those most burdened by inequality — into institutional consequences without requiring Ubuntu to be adopted as a formal policy framework.

These three options are best understood as a spectrum of governability and depth. Where principles dominate without enforceability, Ubuntu is most likely to be absorbed as ethics washing. Where procedural safeguards exist without mechanisms for contestation or remedy, Ubuntu may shape language while leaving distributions

¹ Law relating to the protection of personal data and privacy (Rwanda), 2021, Law No 058/2021 of 13/10/2021.

² EC, Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act), [2024] OJ, L 2024/1689, online: <<https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>>.

of power, risk and harm largely unchanged. Where structural governance is pursued — and procedural tools are treated as enforceable gates rather than as paperwork — the political turn in Ubuntu becomes operational rather than rhetorical.

Policy Recommendations: Translating a Political Reading of Ubuntu into Practice

The recommendations that follow are directed primarily at public-sector governance actors, particularly regulators, public procurement authorities and agencies overseeing high-impact AI deployment, though they also speak to civil society advocates and other institutional actors involved in implementation and oversight.

A political rendering of Ubuntu shifts ethical attention from shared values to the organization of relationships under conditions of unequal power. If persons are constituted through relations — and if AI systems actively reshape those relations by reallocating authority, risk and benefit — then ethics cannot remain at the level of orientation or intent. It must be translated into institutional arrangements that determine who decides, who benefits, who bears harm and how conflicts are addressed. From this perspective, Ubuntu not only prescribes harmony but also demands governance capable of constraining domination, enabling contestation and repairing relational harm. This does not require the formal adoption of Ubuntu as a policy framework; rather, it requires governance mechanisms that deliver these outcomes in practice. The governance moves set out below follow from this logic: They translate Ubuntu's relational commitments into concrete mechanisms that redistribute decision rights, attach obligations to extraction and value creation and establish remedies where AI systems fracture social relations.

Requiring Participatory Governance with Real Decision-Making Power

A political reading of Ubuntu requires that those most affected by AI systems are not treated as stakeholders to be consulted but as agents with

standing in decisions that reshape their lives. In practice, this calls for participatory governance arrangements with real authority in high-impact AI deployments, particularly in public services, biometric identification and surveillance contexts (Young, Katell and Krafft 2019). Public procurement units and sector regulators can operationalize this process by requiring standing oversight bodies as a condition of deployment — bodies composed of affected community representatives, civil society actors and independent technical and legal expertise. These bodies must have defined powers to require modifications, impose conditions, pause rollout or recommend non-deployment where relational harms are likely (Southerland 2023). To sustain legitimacy and avoid capture, such arrangements should include rotating membership with fixed terms, transparent agendas and published records, as well as time-limited approvals subject to mandatory sunset review (Xanthaki 2020). In this way, participation becomes a mechanism for redistributing decision-making authority rather than a consultative add-on, embedding relational accountability directly into AI governance.

Embedding Reciprocity Through Solidarity-Based Data Agreements

The politicization of Ubuntu treats data practices as relations of interdependence rather than as isolated transactions. This brings obligations of reciprocity into view when data and model value are generated from communities. Where data flows out of a community through public-private partnerships, cross-institutional data sharing or model-training arrangements, governance should require solidarity-based data agreements that bind extraction to return (Prainsack et al. 2022). Regulators and public procurement authorities can operationalize this by embedding such obligations in data-sharing and AI procurement contracts (Sanchez-Graells 2024; Mahomed, Loots and Staunton 2022), ensuring that benefits flow back through concrete mechanisms such as shared revenues, funded training and fellowships, service discounts or investment in local data capacity (LeBaron von Baeyer et al. 2024). To prevent these commitments from remaining symbolic, agreements should also require auditability through clear documentation of what data is used and for what purposes, where models are developed or deployed, who captures value and what forms of return are provided (Mahomed, Loots and Staunton 2022). By making reciprocity

a contractual condition rather than a voluntary gesture, solidarity-based agreements alter the terms under which data circulates and address extractive dynamics by redistributing benefits and obligations along the AI value chain (Russo Carroll et al. 2020).

Mandating Restorative Responses for AI Harms

Ubuntu understands harm as more than a technical error or legal violation — it is a rupture in social relations that undermines trust, standing and belonging. When AI systems produce harm in domains such as public services, policing, labour platforms or credit allocation, governance responses that stop at explanation or procedural correction leave this relational damage intact (Wong and Rieder 2025). A governance approach informed by Ubuntu therefore requires restorative mechanisms that prioritize repair alongside accountability. Public agencies and sector regulators can give effect to this by mandating response processes that begin with acknowledgement of harm and extend to timely remedies, including correction of records, compensation where appropriate and removal of wrongful classifications or exclusions.³ Repair must also address relational consequences by restoring access, credibility and social standing, particularly where automated decisions have produced stigma or exclusion (Peterson-Salahuddin 2024). To ensure that harm is not normalized through repetition, serious incidents or recurring patterns should trigger independent review and, where necessary, suspension of system use until failures are addressed.

Supporting Collective Data Stewardship Through an Ubuntu-Informed Data Commons

At the heart of Ubuntu is the recognition that persons are constituted through shared relations, a premise that sits uneasily with data governance models built around individual consent and one-off transactions. Where data generates group-level effects and collective value, treating it as a purely individual asset obscures how risks and benefits

are distributed across communities. An Ubuntu-informed governance approach therefore supports the stewardship of such data as a shared resource, governed through collective arrangements rather than through unilateral access (Sun, Shabaya and Kalema 2025). In practice, this can take the form of data trusts or cooperatives that hold community-relevant data sets under duties of care, maintain public registries of data sets and access decisions, and enforce purpose-binding rules that restrict secondary use without collective approval (Delacroix and Lawrence 2019; Ada Lovelace Institute and UK AI Council 2021). Access becomes conditional on transparency, ongoing oversight and demonstrated reciprocity, including contributions to local capacity or public benefit. Through such arrangements, governance shifts from extractive access toward shared responsibility, embedding relational accountability into the infrastructures that support AI development and deployment.

Conclusion

If Ubuntu enters AI governance primarily as ethical language, it risks reinforcing the governance gaps it is often invoked to address. High-level principles and procedural safeguards remain necessary as baseline governance tools but where they operate without enforceable consequences, they allow ethical commitments to circulate while leaving relations of power, extraction and harm largely intact. In such settings, Ubuntu functions less as a corrective than as a legitimizing vocabulary.

The question is therefore not whether Ubuntu should inform AI governance but whether it will have institutional consequence. A political reading of Ubuntu shifts attention from value affirmation to the organization of decision-making authority, the distribution of benefit and risk, and the availability of contestation and remedy.

The governance moves outlined here do not require Ubuntu's adoption as a formal policy framework. They instead translate its relational demands into concrete mechanisms that constrain domination, enable participation and support relational repair. Whether Ubuntu matters in practice depends not on rhetoric but on whether governance systems are willing to give it force.

³ EC, Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), [2016] OJ, L 119/1, online: <<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32016R0679>>.

Works Cited

- Ada Lovelace Institute and UK AI Council. 2021. *Exploring legal mechanisms for data stewardship*. Working group final report. March. London, UK: Ada Lovelace Institute and UK AI Council. www.adalovelaceinstitute.org/report/legal-mechanisms-data-stewardship/.
- Adams, Rachel. 2021. "Can artificial intelligence be decolonized?" *Interdisciplinary Science Reviews* 46 (1–2): 176–97. <https://doi.org/10.1080/03080188.2020.1840225>.
- African Union Commission. 2024. *Continental Artificial Intelligence Strategy: Harnessing AI for Africa's Development and Prosperity*. July. https://au.int/sites/default/files/documents/44004-doc-EN-Continental_AI_Strategy_July_2024.pdf.
- Biddle, Justin B., John P. Nelson and Olajide E. Olugbade. 2024. "How Can We Know if You are Serious? Ethics Washing, Symbolic Ethics Offices, and the Responsible Design of AI Systems." *Canadian Journal of Philosophy* 54 (4): 329–45. <https://doi.org/10.1017/can.2025.9>.
- Bietti, Elettra. 2021. "From Ethics Washing to Ethics Bashing: A Moral Philosophy View on Tech Ethics." *Journal of Social Computing* 2 (3): 266–83. <https://doi.org/10.23919/JSC.2021.0031>.
- Birhane, Abeba. 2021. "Algorithmic injustice: a relational ethics approach." *Patterns* 2 (2): 100205. <https://doi.org/10.1016/j.patter.2021.100205>.
- Couldry, Nick and Ulises A. Mejias. 2019. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford, CA: Stanford University Press.
- Crawford, Kate. 2021. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven, CT: Yale University Press.
- de Vries, Jantina. 2024. "Centering Africa as context and driver for Global Health Ethics: incompleteness, conviviality and the limits of Ubuntu." *Wellcome Open Research* 9: 371. <https://doi.org/10.12688/wellcomeopenres.22508.1>.
- Delacroix, Sylvie and Neil D. Lawrence. 2019. "Bottom-up data Trusts: disturbing the 'one size fits all' approach to data governance." *International Data Privacy Law* 9 (4): 236–52. <https://doi.org/10.1093/idpl/ipz014>.
- Eke, Damian and George Ogoh. 2022. "Forgotten African AI Narratives and the future of AI in Africa." *International Review of Information Ethics* 31 (1). <https://doi.org/10.29173/irie482>.
- Ewuoso, Cornelius. 2021. "An African Relational Approach to Healthcare and Big Data Challenges." *Science and Engineering Ethics* 27, 34. <https://doi.org/10.1007/s11948-021-00313-w>.
- Ferlito, Brandon, Seppe Segers, Michiel De Proost and Heidi Mertes. 2024. "Responsibility Gap(s) Due to the Introduction of AI in Healthcare: An Ubuntu-Inspired Approach." *Science and Engineering Ethics* 30, 34. <https://doi.org/10.1007/s11948-024-00501-4>.
- Gwagwa, Arthur, Emre Kazim and Airlie Hilliard. 2022. "The role of the African value of Ubuntu in global AI inclusion discourse: A normative ethics perspective." *Patterns* 3 (4): 100462. <https://doi.org/10.1016/j.patter.2022.100462>.
- Hagendorff, Thilo. 2020. "The Ethics of AI Ethics: An Evaluation of Guidelines." *Minds and Machines* 30 (1): 99–120. <https://doi.org/10.1007/s11023-020-09517-8>.
- Hao, Karen. 2019. "In 2020, let's stop AI ethics-washing and actually do something." *MIT Technology Review*, December 27. www.technologyreview.com/2019/12/27/57/ai-ethics-washing-time-to-act/.
- Jobin, Anna, Marcello Lenca and Effy Vayena. 2019. "The global landscape of AI ethics guidelines." *Nature Machine Intelligence* 1 (9): 389–99. <https://doi.org/10.1038/s42256-019-0088-2>.
- LeBaron von Baeyer, Sarah, Rebecca Crocker, Rindra Rakotoarivony, Jean Freddy Ranaivoarisoa, Germain Jules Spiral, Stephane Castel, Andrew Farnum et al. 2024. "Why community consultation matters in genomic research benefit-sharing models." *Genome Research* 34 (1): 1–6. <https://doi.org/10.1101/gr.278308.123>.
- Mahomed, Safia, Glaudina Loots and Ciara Staunton. 2022. "The role of Data Transfer Agreements in ethically managing data sharing for research in South Africa." *South African Journal of Bioethics and Law* 15 (1): 26–30. <https://doi.org/10.7196/SAJBL.2022.v15i1.807>.
- Mensah, Jahaziel Osei and Aimee Van Wynsberghe. 2025. "Where are the missing values: an exploration of the need to incorporate Ubuntu values into African AI policy." *AI and Ethics* 5: 4365–75. <https://doi.org/10.1007/s43681-025-00746-0>.
- Mhlambi, Sábêlo and Simona Tiribelli. 2023. "Decolonizing AI Ethics: Relational Autonomy as a Means to Counter AI Harms." *Topoi* 42 (3): 867–80. <https://doi.org/10.1007/s11245-022-09874-2>.
- Mohamed, Shakir, Marie-Therese Png and William Isaac. 2020. "Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence." *Philosophy & Technology* 33 (4): 659–84. <https://doi.org/10.1007/s13347-020-00405-8>.
- Molefe, Motsamai. 2025. *Ubuntu Ethics: Human Dignity, Moral Perfectionism, and Needs*. Routledge Studies in African Philosophy Series. Abingdon, UK: Routledge.
- Muya, Catherine. 2021. "What Kenya needs to do to get it right on Huduma Namba." Opinion, *The Star*, October 19. www.the-star.co.ke/opinion/star-blogs/2021-10-24-what-kenya-needs-to-do-to-get-it-right-on-huduma-namba.
- National Institute of Standards and Technology. 2023. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. January. Gaithersburg, MD: National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>.
- Organisation for Economic Co-operation and Development. 2025. *Recommendation of the Council on Artificial Intelligence*. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>.

- Odero, Brenda, David Nderitu and Gabrielle Samuel. 2024. "The Ubuntu Way: Ensuring Ethical AI Integration in Health Research." *Wellcome Open Research* 9: 625. <https://doi.org/10.12688/wellcomeopenres.23021.1>.
- Peterson-Salahuddin, Chelsea. 2024. "Repairing the harm: Toward an algorithmic reparations approach to hate speech content moderation." *Big Data & Society* 11 (2). <https://doi.org/10.1177/20539517241245333>.
- Prainsack, Barbara, Selim El-Sayed, Nikolaus Forgó, Łukasz Szoszkiewicz and Philipp Baumer. 2022. "Data solidarity: a blueprint for governing health futures." *The Lancet Digital Health* 4 (11): e773–74. [https://doi.org/10.1016/S2589-7500\(22\)00189-3](https://doi.org/10.1016/S2589-7500(22)00189-3).
- Rességuier, Anaïs and Rowena Rodrigues. 2020. "AI ethics should not remain toothless! A call to bring back the teeth of ethics." *Big Data & Society* 7 (2): 2053951720942541. <https://doi.org/10.1177/2053951720942541>.
- Reviglio, Urbano and Rogers Alunge. 2020. "'I Am Datafied Because We Are Datafied': an Ubuntu Perspective on (Relational) Privacy." *Philosophy & Technology* 33 (4): 595–612. <https://doi.org/10.1007/s13347-020-00407-6>.
- Russo Carroll, Stephanie, Ibrahim Garba, Oscar L. Figueroa-Rodríguez, Jarita Holbrook, Raymond Lovett, Simeon Materechera, Mark Parsons et al. 2020. "The CARE Principles for Indigenous Data Governance." *Data Science Journal* 19: 43, 1–12. <https://doi.org/10.5334/dsj-2020-043>.
- Sanchez-Graells, Albert. 2024. "Responsibly Buying Artificial Intelligence: A 'Regulatory Hallucination.'" *Current Legal Problems* 77 (1): 81–126. <https://doi.org/10.1093/clp/cuae003>.
- Segun, Samuel T. 2021. "Critically engaging the ethics of AI for a global audience." *Ethics and Information Technology* 23 (2): 99–105. <https://doi.org/10.1007/s10676-020-09570-y>.
- . 2024. "Are certain African ethical values at risk from artificial intelligence?" *Data & Policy* 6: e68. <https://doi.org/10.1017/dap.2024.64>.
- Southerland, Vincent M. 2023. "The Master's Tools and a Mission: Using Community Control and Oversight Laws to Resist and Abolish Police Surveillance Technologies." *UCLA Law Review* 70 (2): 1–90. www.uclalawreview.org/the-masters-tools-and-a-mission-using-community-control-and-oversight-laws-to-resist-and-abolish-police-surveillance-technologies/.
- Sun, Simon (Chieh-Jan), Marie Muhadia Shabaya and Nai Lee Kalema. 2025. "Fostering African Data Commons: Embracing the Philosophy of Ubuntu." In *Oxford Intersections: AI in Society*, edited by Philipp Hacker. Oxford, UK: Oxford University Press. <https://doi.org/10.1093/9780198945215.003.0044>.
- United Nations Educational, Scientific and Cultural Organization. 2021. *Recommendation on the ethics of artificial intelligence*. <https://unesdoc.unesco.org/ark:/48223/pf0000380455/PDF/380455eng.pdf.multi>.
- van Maanen, Gijs. 2022. "AI Ethics, Ethics Washing, and the Need to Politicize Data Ethics." *Digital Society* 1: 9. <https://doi.org/10.1007/s44206-022-00013-3>.
- van Norren, Doreen Eva. 2023. "The ethics of artificial intelligence, UNESCO and the African Ubuntu perspective." *Journal of Information, Communication and Ethics in Society* 21 (1): 112–28. <https://doi.org/10.1108/JICES-04-2022-0037>.
- Wong, Pak-Hang and Gernot Rieder. 2025. "After Harm: A Plea for Moral Repair after Algorithms Have Failed." *Science and Engineering Ethics* 31 (5): 26. <https://doi.org/10.1007/s11948-025-00555-y>.
- Xanthaki, Helen. 2020. "Sunset Clauses: A Contribution to Legislative Quality." In *Time, Law, and Change: An Interdisciplinary Study*, edited by Sofia Ranchordá and Yaniv Roznai, 209–26. Oxford, UK: Hart. <https://doi.org/10.5040/9781509930968.ch-009>.
- Yeung, Karen. 2018. "Algorithmic regulation: A critical interrogation." *Regulation & Governance* 12 (4): 505–23. <https://doi.org/10.1111/rego.12158>.
- Yilma, Kinfe. 2025. "Ethics of AI in Africa: Interrogating the role of Ubuntu and AI governance initiatives." *Ethics and Information Technology* 27 (2): 1–14. <https://doi.org/10.1007/s10676-025-09834-5>.
- Young, Meg, Michael Katell and P. M. Krafft. 2019. "Municipal surveillance regulation and algorithmic accountability." *Big Data & Society* 6 (2). <https://doi.org/10.1177/2053951719868492>.

About CIGI

The Centre for International Governance Innovation (CIGI) is an independent, non-partisan think tank whose peer-reviewed research and trusted analysis influence policy makers to innovate. Our global network of multidisciplinary researchers and strategic partnerships provide policy solutions for the digital era with one goal: to improve people's lives everywhere. Headquartered in Waterloo, Canada, CIGI has received support from the Government of Canada, the Government of Ontario and founder Jim Balsillie.

À propos du CIGI

Le Centre pour l'innovation dans la gouvernance internationale (CIGI) est un groupe de réflexion indépendant et non partisan dont les recherches évaluées par des pairs et les analyses fiables incitent les décideurs à innover. Grâce à son réseau mondial de chercheurs pluridisciplinaires et de partenariats stratégiques, le CIGI offre des solutions politiques adaptées à l'ère numérique dans le seul but d'améliorer la vie des gens du monde entier. Le CIGI, dont le siège se trouve à Waterloo, au Canada, bénéficie du soutien du gouvernement du Canada, du gouvernement de l'Ontario et de son fondateur, Jim Balsillie.

Credits

Director, Programs **Dianna H. English**
Senior Program Manager **Ifeoluwa Olorunnipa**
Publications Editor **Christine Robertson**
Graphic Designer **Sepideh Shomali**

Copyright © 2026 by the Centre for International Governance Innovation

The opinions expressed in this publication are those of the author and do not necessarily reflect the views of the Centre for International Governance Innovation or its Board of Directors.

For publications enquiries, please contact publications@cigionline.org.



The text of this work is licensed under CC BY 4.0. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

For reuse or distribution, please include this copyright notice. This work may contain content (including but not limited to graphics, charts and photographs) used or reproduced under licence or with permission from third parties. Permission to reproduce this content must be obtained from third parties directly.

Centre for International Governance Innovation and CIGI are registered trademarks.

67 Erb Street West
Waterloo, ON, Canada N2L 6C2
cigionline.org

