

Digital Policy Hub – Working Paper

Open Weights, Hidden Dependencies: Governing Open-Weight AI in Africa

Amal Hussein

Winter term – 2025–2026 cohort

About the Hub

The Digital Policy Hub at CIGI is a collaborative space for emerging scholars and innovative thinkers from the social, natural and applied sciences. It provides opportunities for undergraduate and graduate students and post-doctoral and visiting fellows to share and develop research on the rapid evolution and governance of transformative technologies. The Hub is founded on transdisciplinary approaches that seek to increase understanding of the socio-economic and technological impacts of digitalization and improve the quality and relevance of related research. Core research areas include data, economy and society; artificial intelligence; outer space; digitalization, security and democracy; and the environment and natural resources.

The Digital Policy Hub working papers are the product of research related to the Hub's identified themes prepared by participants during their fellowship.

Partners

Thank you to Mitacs for its partnership and support of Digital Policy Hub fellows through the Accelerate program. We would also like to acknowledge the many universities, governments and private sector partners for their involvement allowing CIGI to offer this holistic research environment.



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.

About the Author

Amal Hussein is a Digital Policy Hub master's fellow and an M.Sc. candidate in systems science and engineering with a concentration in interdisciplinary artificial intelligence (AI) at the University of Ottawa. She contributes research on AI ethics and governance with the Canadian Robotics and Artificial Intelligence for Engineering Design Laboratory. A multidisciplinary designer with a background in systems design and applied AI, Amal explores how global AI ethics principles align with African political, economic and governance contexts. Her work spans public interest technology, civic tech innovation and responsible systems design, and is informed by her prior experience in federal digital service delivery.

Copyright © 2026 by Amal Hussein

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.

Centre for International Governance Innovation and CIGI are registered trademarks.

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

Open Weights, Hidden Dependencies: Governing Open-Weight AI in Africa

Amal Hussein

Bottom Line Up Front

Open-weight artificial intelligence (AI) foundation models are systems whose trained parameters are publicly released for download and local deployment. They are increasingly capable and widely available, and represent a lower-cost pathway to technological sovereignty for governments globally. As open-weight models become more accessible, the barrier to adoption is no longer technical readiness. The absence of operational governance frameworks — which specify accountability structures, pre-deployment requirements and institutional responsibility for system behaviour — remains the primary obstacle to defensible adoption. African public institutions face this challenge under conditions of uneven institutional capacity and significant infrastructure gaps. Tiered policy recommendations addressing national procurement, update governance and regional coordination can enable safe open-weight adoption without waiting for comprehensive governance infrastructure that may remain out of reach.

Key Points

- **Open-weight models already exist and are freely available:** Whether African institutions can deploy them in ways that build institutional capacity and advance technological sovereignty is the question that matters.
- **Unlike closed systems, open-weight models shift governance responsibility to the point of deployment:** Because weights can be freely downloaded, locally deployed, customized and modified without developer oversight, the focus of effective AI governance moves from model development to the far more varied context of adoption. This creates both opportunity and accountability.
- **The governance gap is addressable:** 44 of 54 African countries have data protection laws, and 38 have fully operational data protection authorities. These frameworks can anchor the accountability structures and security controls that open-weight deployment requires without waiting for AI-specific legislation.

Recommendations

- **Establish minimum deployment baselines:** African governments should require pre-deployment risk classification, input-output logging and documented incident response procedures for any open-weight model used in public service delivery. These controls are low-cost and auditable, and anchor directly to existing data protection frameworks.
- **Govern model updates at the point of deployment:** Governments should require designated update custodians, documented version control and re-evaluation procedures when model weights are modified or fine-tuned by contractors. Accountability must follow the model through its full life cycle, reaching the actors most likely to alter it after initial deployment.
- **Invest in regional governance and infrastructure coordination:** African Union member states should establish shared security evaluation standards and cross-border audit frameworks. Pooled compute infrastructure and a continental mechanism for data governance coordination would reduce regulatory fragmentation and give African institutions the capacity to run, adapt and govern open-weight models on terms that they control.

Introduction

A recurring debate in artificial intelligence (AI) governance discourse frames Africa's choices around whether governments should invest limited resources in closed, proprietary AI systems or open-weight alternatives. The framing reflects substantive concerns about participation in a rapidly concentrating global AI economy. It also obscures a more tractable opportunity. Open-weight foundation models already exist, and are publicly available and increasingly capable. The question worth asking is whether African institutions can deploy and govern these models in ways that build the AI ecosystem that the continent needs.

An open-weight model is one whose trained parameters are publicly released for download and local deployment. This is technically distinct from open source in the traditional software sense. The Organisation for Economic Co-operation and Development clarifies that calling AI models open source is often misleading when only weights are released, while training data, code and documentation remain proprietary (Organisation for Economic Co-operation and Development 2025). What the release of weights enables is local deployment without ongoing cloud dependence, adaptation to local languages and use cases, and institutional control over how and whether the model is updated. Whether African governments integrate open-weight models into their AI strategies may prove to be one of the more consequential decisions shaping their long-term position in the global AI economy (Surman 2025).

The properties of open-weight models create an opportunity that extends beyond cost savings. Frontier AI companies have responded to sovereignty concerns by building local data centres but data residency does not equal sovereignty. Under the US CLOUD (Clarifying Lawful Overseas Use of Data) Act, American courts can compel access to data held by US-based companies regardless of where it is physically stored (Dale 2025). Vendor control over model updates and restrictions on fine tuning compound the dependency further (Widder, Whittaker and West 2024; Unger and McLean 2025). Closed models also run at roughly six times the price of open alternatives (Nagle and Yue 2025).

Open-weight models create conditions under which African institutions can take meaningful ownership of their AI systems, but that relationship between open-weight availability and African AI sovereignty is not automatic. Without compute infrastructure that African institutions can access and control, those institutions remain dependent on externally controlled access. Individual institutional ownership of that infrastructure is out of reach for most: only five percent of African AI practitioners currently have access to adequate computational resources, making regional infrastructure sharing the more viable path to genuine compute sovereignty (Adan 2025). Governance is what converts open-weight availability into open-weight opportunity.

Africa's AI governance agenda has advanced significantly at the strategic level. The 2024 African Union (AU) Continental AI Strategy, the 2025 Africa Declaration on Artificial Intelligence and the establishment of the Africa AI Council collectively signal continental commitment to inclusive and sovereign AI development (Munga and Quansah 2025). The risk is not a lack of ambition but divergence between declaration and operational practice, and a narrowing window of influence as the major global AI governance convenings that elevated African priorities in 2024 and 2025 shift away from the

continent in 2026, with the Group of Twenty presidency moving to the United States and no major Global South-hosted AI governance summit currently scheduled (Vithani 2026). Moving from strategy to implementation requires concrete mechanisms (Munga and Quansah 2025). Open-weight governance is one such mechanism.

The central question is: what minimum security controls and governance steps enable African public institutions to adopt open-weight foundation models safely for service delivery, given uneven institutional capacity and infrastructure gaps across the continent?

From Performance Gap to Governance Barrier

The narrowing performance gap between open-weight and closed models has removed a key technical justification for defaulting to proprietary systems. The Stanford Institute for Human-Centered Artificial Intelligence reports that by February 2025, the performance gap between leading closed-weight and open-weight models on the chatbot arena leaderboard had narrowed to 1.70 percent, down from 8.04 percent in January 2024 (Maslej et al. 2025). Open-weight models now represent over 55 percent of all commercially available foundation models worldwide (Organisation for Economic Co-operation and Development 2025).

Open-weight systems raise distinct challenges that general AI governance frameworks do not fully address. Weights can be modified after distribution, safety constraints can be removed and update governance can be fragmented across agencies or contractors (Casper et al. 2025). A model that passes initial evaluation may be modified after deployment to remove alignment constraints, introduce biases or create security vulnerabilities. For public institutions without continuous technical monitoring capacity, this potential for modification creates ongoing exposure that differs qualitatively from the risks of closed systems. These risks are not unique to open-weight systems, but the absence of a vendor managing security patches and mitigations means that the deploying institution bears full responsibility for monitoring and response. Prompt injection, insecure tool use and data leakage are among the most consequential threats in deployed generative AI systems (Open Worldwide Application Security Project 2024), and each carries amplified risks in public-sector contexts where model outputs may influence decisions about benefits, health information or public safety.

Organizational dynamics compound the technical picture. Research on public sector AI adoption shows that institutional hesitancy is shaped by accountability concerns, ethical uncertainty and unclear regulatory expectations, each of which raises the threshold for action and intensifies as systems become more exposed to external users (Haesevoets, Verschuere and Roets 2025). For open-weight deployments specifically, these dynamics are amplified: without a vendor to share accountability when something goes wrong, the full burden of justifying deployment falls on the institution itself. Open-weight models are sufficiently capable for the tasks that public institutions most need them to perform. Information delivery, document processing and citizen-facing service assistance are well within the capability range of currently available open-weight systems, including those that can be fine-tuned for local languages and

contexts. What institutions lack is governance clarity: the documented controls, defined accountability and procedural requirements that make public-sector deployment defensible. This working paper proposes a three-tier governance framework covering baseline deployment controls, update governance protocols and regional coordination mechanisms, with each tier designed to be implementable under the institutional and infrastructure conditions that currently exist across the continent.

African Contexts: Infrastructure, Sovereignty and Regional Opportunity

Africa's AI governance architecture has developed rapidly at the strategic level while the operational level remains underdeveloped. The continent hosts less than one percent of global data centres and three percent of global AI talent (Munga and Quansah 2025). These conditions define the starting point from which governance frameworks must be designed if they are to be usable at all.

The gap between continental frameworks and institutional practice sits at the heart of the adoption problem. The AU Continental AI Strategy's implementation timeline runs from 2025 to 2030, with phase one focused on establishing governance structures and national strategies (AU 2024). For institutions facing deployment decisions now, that timeline offers limited operational guidance. What minimum governance at the institutional level can do is reduce fragmentation by giving individual ministries, agencies and public bodies something concrete to implement while broader coordination develops. Forty-four of 54 African countries have enacted data protection laws, and 38 have fully operational data protection authorities (Buza 2026). These frameworks provide a ready legal foundation for AI governance because they already establish the accountability structures, documentation requirements and breach notification obligations that open-weight deployment controls require. Where AI-specific regulation is absent, data protection law can anchor the security controls, documentation requirements and accountability mechanisms needed to address the specific risks of open-weight deployment — namely weight modification after distribution, fragmented update governance and unclear institutional liability — without waiting for new legislation. Governance controls of the kind proposed here — including risk classification, logging, incident response and update custodian — are auditable in principle but require oversight mechanisms beyond the deploying institution itself to function as intended. Civil society scrutiny, parliamentary oversight and the cross-border enforcement cooperation that African data protection authorities are increasingly exercising (ibid.) are the conditions under which minimum viable governance produces genuine accountability rather than documented compliance.

Open-weight models are well-suited to African public-sector conditions in ways that closed systems are not. Local deployment means sensitive public data need not leave national jurisdiction for processing. The availability of weights for inspection and evaluation supports the kind of independent audit and accountability that public institutions require. Fine-tuning without returning to a vendor enables adaptation to local languages and contexts that proprietary systems often cannot accommodate without significant additional cost or negotiation. These properties make open-weight models a viable entry point into AI development that is genuinely African-led rather than externally licensed.

The opportunity, however, depends on infrastructure that does not yet exist at regional scale. Open-weight AI only becomes meaningful when African institutions can actually run, fine-tune and evaluate models on infrastructure that they control. The governance framework proposed in this working paper assumes that access to such infrastructure exists; it does not address the upstream challenge of building locally representative data sets at population scale, which represents a distinct investment requirement that regional governance coordination alone cannot resolve. At present, many institutions rely on short-term cloud credits or externally managed access, which reproduces the dependency that open-weight adoption is meant to reduce. Regional compute infrastructure — shared across universities, public agencies and research institutions because individual institutional ownership remains out of reach for most — would allow these actors to work with open-weight models on terms that build lasting capacity (Adan 2025). A continental framework for data governance coordination would address the legal complexity that currently limits cross-border data collaboration, even among institutions that share development goals and political will. These investments are what make open-weight governance durable and give African institutions the capacity to govern AI on their own terms.

Governance Analysis

Effective governance for open-weight deployment in African public-sector contexts requires controls that operate at three distinct levels of institutional action:

- At the level of individual deployments, baseline controls establish the minimum conditions for defensible adoption.
- At the level of ongoing institutional responsibility, update governance addresses the accountability gap that opens when locally deployed models are modified after initial evaluation.
- At the regional level, coordination mechanisms make the sovereignty benefits of open-weight models practically achievable rather than theoretically available.

Tier 1: Baseline Controls for All Open-Weight Public Deployments

Baseline controls must be achievable by institutions with varying levels of technical capacity, as well as generating auditable documentation and addressing the risks most salient to public-sector deployments: data exposure, accountability gaps and undesirable model behaviour. Three baseline controls meet these criteria.

The first is pre-deployment risk classification. Before deploying any open-weight model in a service delivery context, institutions should classify the deployment along three dimensions: the sensitivity of data processed (personal, sensitive or public); the nature of decision influence (meaning whether the model output provides general information, informs a human decision or directly determines an outcome affecting a citizen); and the degree of modification from the base model weights (unmodified, fine-tuned or substantially modified). This classification determines which additional controls apply and provides the documentation basis for accountability. Risk classification can be completed using standard data governance training and

aligns directly with data protection frameworks already in place across the African continent, making it accessible to institutions without dedicated AI teams.

The second is input-output logging with defined retention and access controls. Logging model inputs and outputs serves multiple governance functions: it enables incident investigation, supports audits of decision influence and provides the evidence base for monitoring model behaviour over time (Open Worldwide Application Security Project 2024; National Institute of Standards and Technology 2023). Logging must be paired with access controls and defined retention periods under data protection frameworks. For deployments processing sensitive personal data, pseudonymization of logged content should be required.

The third is documented incident response procedures. The National Institute of Standards and Technology's AI Risk Management Framework identifies incident response as a core life cycle governance function (National Institute of Standards and Technology 2023). For open-weight deployments, incident response must address model-specific scenarios including prompt injection attacks, unexpected model outputs and post-deployment weight modifications. The Open Worldwide Application Security Project Top 10 for Large Language Model Applications provides the threat taxonomy that incident response procedures should address (Open Worldwide Application Security Project 2024). Procedures should specify response authorities, escalation paths and communication requirements.

Tier 2: Update Governance for Locally Deployed Models

The most distinctive governance challenge of open-weight deployment is update responsibility. When a vendor manages a closed model, the institution's governance obligation is primarily to evaluate the system at deployment and monitor its behaviour in use. When an institution deploys an open-weight model locally, it becomes responsible for all subsequent updates including security patches, fine-tuning and weight modifications by contractors. Without update governance, a model that is evaluated and approved at deployment can be materially altered without triggering re-evaluation.

Institutions deploying open-weight models should designate an update custodian: a role or office responsible for authorizing and documenting model updates. Version control requirements should specify that each update to model weights must be documented with a change description, authorization record and risk classification review. Updates that materially change model capabilities or the risk profile of the deployment should trigger re-evaluation under tier 1 baseline controls. For deployments classified as higher-risk under tier 1, including those processing sensitive health data or influencing benefit eligibility decisions, re-evaluation should be required for any weight modification regardless of stated scope.

Procurement clauses should operationalize update governance at the contract stage. When institutions procure deployment support, customization services or fine-tuning from contractors, contracts should specify update documentation requirements, notification obligations when model weights are modified and data handling conditions during fine-tuning. These clauses extend governance obligations to the third-party actors most likely to modify model weights after initial deployment.

Tier 3: Regional Coordination and Infrastructure

The cost of comprehensive AI security governance exceeds what most public institutions can sustain independently. Regional coordination can distribute these costs while building the shared infrastructure that makes open-weight adoption sustainable and institutionally controlled, rather than contingent on short-term cloud credits or externally managed compute access.

The first mechanism is shared security evaluation standards. The AU Continental AI Strategy calls for mechanisms for sharing AI expertise and best practices between member countries (AU 2024). Concretely, this should include a shared library of evaluation criteria and test prompts for common open-weight deployment scenarios in, for example, health care information delivery, agricultural extension services and social benefit administration. Shared standards enable institutions to conduct rigorous evaluations without developing criteria independently and support mutual recognition of evaluations across jurisdictions (Buza 2026).

The second mechanism is pooled security audit capacity. A regional AI security audit pool housed within an existing AU institution or established as a new body could provide audit services to member states on a shared-cost basis. This model has precedent in other domains of technical regulatory cooperation and would create a regional repository of deployment experience that individual institutions cannot accumulate alone.

The third mechanism is shared regional compute infrastructure paired with continental data governance coordination. Regional investment in shared compute capacity — whether through development bank financing or pooled government procurement — would allow universities, public agencies and research institutions to run, fine-tune and evaluate open-weight models on infrastructure under regional control. The governance framework underpinning such infrastructure would need to include explicit data sovereignty protections ensuring that each nation retains legal control over its data regardless of the physical location of the shared facility, with intergovernmental agreements establishing clear boundaries on what host countries can access. Regional compute infrastructure that lacks these protections risks reproducing the same sovereignty concerns it is designed to resolve (Adan 2025). A continental framework for data governance coordination would complement this investment by reducing the legal fragmentation that currently limits what institutions can do together.

Recommendations

- **Recommendation 1: Establish minimum deployment baselines.** African governments should require pre-deployment risk classification, input-output logging and documented incident response procedures for any open-weight model deployed in public service delivery. Risk classification should assess data sensitivity, the nature of decision influence and the ethical implications of deployment. Logging requirements should be paired with retention limits and access controls under existing data protection law. Incident response procedures should address the threat scenarios most salient to generative AI deployments, including prompt injection and post-deployment weight modification. These controls do not require new legislation and can be implemented by institutions at varying levels of technical capacity.

- **Recommendation 2: Adopt update governance protocols for locally deployed models.** Governments should require designated update custodians, documented version control and re-evaluation procedures when model weights are modified or fine-tuned. Procurement contracts with vendors and contractors should specify update documentation requirements and notification obligations. For higher-risk deployments processing sensitive personal data or influencing consequential decisions, any weight modification should trigger re-evaluation under the baseline controls in recommendation 1. These protocols extend accountability to the point of greatest post-deployment risk.
- **Recommendation 3: Invest in regional governance and infrastructure coordination.** AU member states should establish shared security evaluation standards and a regional AI security audit mechanism that member states can access on a shared-cost basis. In parallel, regional investment in shared compute infrastructure and a continental framework for data governance coordination would address the infrastructure and legal conditions that currently limit what African institutions can do with open-weight AI. These investments make the sovereignty benefits of open-weight models practically achievable rather than theoretically available.

Conclusion

Open-weight models exist, and are capable and freely accessible. The institutions and infrastructure needed to deploy them on African terms are buildable, and the governance frameworks that make deployment safe and accountable are the foundation on which that building begins.

Baseline controls are achievable across varying institutional contexts and anchor to data protection frameworks that already exist in 38 of 54 African countries. Update governance protocols could address the distinctive accountability challenge of locally deployed models, ensuring that institutions remain responsible for what their systems do after initial deployment. Regional coordination mechanisms could distribute governance costs and point toward the shared compute infrastructure and data governance coordination bodies that would make open-weight sovereignty durable.

2026 is a consolidation year for African-led AI governance: a period to institutionalize gains rather than announce new frameworks (Vithani 2026). These recommendations do not require new legislation, continental bodies or financing mechanisms to begin; they require institutions to make concrete decisions about who is responsible for what, to document those decisions and to share what they learn. That is minimum viable governance. It is also the practice through which African institutions develop the capacity to govern AI on their own terms.

Acknowledgements

I am grateful to Dr. Jason Millar, my supervisor at the University of Ottawa, for his encouragement, guidance and support as this work developed. I also sincerely thank Duncan Cass-Beggs, my CIGI mentor, for his thoughtful feedback and perspective, and Anna Legault, my Hub peer reviewer, for the care and attention she brought to reviewing this paper. I am grateful to the 2025–2026 Digital Policy Hub cohort for a year of meaningful exchange and shared learning. Finally, I thank Dianna English and Meaghan Dietrich for their continued support and generosity throughout my Hub fellowship.

Works Cited

- Adan, Sumaya Nur. 2025. "Strengthening Africa's AI Ecosystem Through Open-Source Innovation and Policy Alignment." Global Center on AI Governance, April 13. www.globalcenter.ai/research/strengthening-africa-s-ai-ecosystem-through-open-source-innovation-and-policy-alignment.
- AU. 2024. *Continental Artificial Intelligence Strategy for Africa: 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.
- Buza, Maria. 2026. "The Year of the Teeth: Data Protection in Africa Roundup, 2025, Projections for 2026." Digital Policy Alert, January 11. <https://digitalpolicyalert.org/blog/data-protection-in-africa-roundup>.
- Casper, Stephen, Kyle O'Brien, Shayne Longpre, Elizabeth Seger, Kevin Klyman, Rishi Bommasani, Aniruddha Nrusimha et al. 2025. "Open Technical Problems in Open-Weight AI Model Risk Management." SSRN, October 26. <http://dx.doi.org/10.2139/ssrn.5705186>.
- Dale, Robert. 2025. "Sovereign AI in 2025." *Natural Language Processing* 31 (5): 1312–21. <https://doi.org/10.1017/nlp.2025.10007>.
- Haesevoets, Tessa, Bram Verschuere and Arne Roets. 2025. "AI adoption in public administration: Perspectives of public sector managers and public sector non-managerial employees." *Government Information Quarterly* 42 (2): 102029. <https://doi.org/10.1016/j.giq.2025.102029>.
- Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Njenga Kariuki, Emily Capstick et al. 2025. *Artificial Intelligence Index Report 2025*. Preprint, arXiv. <https://doi.org/10.48550/arXiv.2504.07139>.
- Munga, Jane and Sylvester Quansah. 2025. "Understanding Africa's AI Governance Landscape: Insights From Policy Practice and Dialogue." Carnegie Endowment for International Peace, September 11. <https://carnegieendowment.org/posts/2025/09/understanding-africas-ai-governance-landscape-insights-from-policy-practice-and-dialogue>.
- Nagle, Frank and Daniel Yue. 2025. "The Latent Role of Open Models in the AI Economy." Georgia Tech Scheller College of Business Research Paper No. 5767103. SSRN, November 18. <http://dx.doi.org/10.2139/ssrn.5767103>.
- National Institute of Standards and Technology. 2023. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. January. <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>.
- Open Worldwide Application Security Project. 2024. *OWASP Top 10 for Large Language Model Applications 2025*. Version 2025. November 18. <https://genai.owasp.org/resource/owasp-top-10-for-llm-applications-2025/>.

- Organisation for Economic Co-operation and Development. 2025. *AI Openness: A Primer for Policymakers*. OECD AI Paper No. 44. August. www.oecd.org/content/dam/oecd/en/publications/reports/2025/08/ai-openness_958d292b/02f73362-en.pdf.
- Surman, Mark. 2025. "Open source could pop the AI bubble – and soon." *Financial Times*, December 17. www.ft.com/content/353adf48-a5d6-4ea0-b5d7-9a94b141213a.
- Unger, Noam and Madeleine McLean. 2025. *An Open Door: AI Innovation in the Global South amid Geostrategic Competition*. Center for Strategic and International Studies, August 13. www.csis.org/analysis/open-door-ai-innovation-global-south-amid-geostrategic-competition.
- Vithani, Priya. 2026. "Seizing the Watershed Moment: Africa's AI Strategy." CIGI Policy Brief No. 224. February 12. Waterloo, ON: Centre for International Governance Innovation. www.cigionline.org/publications/seizing-the-watershed-moment-africas-ai-strategy/.
- Widder, David Gray, Meredith Whittaker and Sarah Myers West. 2024. "Why 'open' AI systems are actually closed, and why this matters." *Nature* 635: 827–33. <https://doi.org/10.1038/s41586-024-08141-1>.