

Document Version

Final published version

Licence

CC BY-NC-ND

Citation (APA)

Spencer, E. J., Singh, J., Moodley, K., van den Hoven, J., Hazelhorst, J., Gommers, D., & van Genderen, M. E. (2026). Why responsible AI needs regional networks in low-resource health systems. *npj Digital Medicine*, 9(1), Article 725. <https://doi.org/10.1038/s41746-026-03208-x>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

<https://doi.org/10.1038/s41746-026-03208-x>

Why responsible AI needs regional networks in low-resource health systems



Emma-Jane Spencer^{1,2}, Jerome Singh³, Keymanthri Moodley⁴, Jeroen van den Hoven⁵, Jip Hazelhorst², Diederik Gommers^{1,2} & Michel E. van Genderen^{1,2,6} ✉

Responsible clinical AI requires regionally anchored networks tailored to local contexts. We propose a primary-care-first approach to prioritize local governance, ensuring decisions are context-driven; data sovereignty, to build trust and protect privacy; and community-driven infrastructure, ensuring sustainability. With capacity building, these networks support equitable AI by aligning with local needs, overcoming barriers to implementation, and enabling collaboration across regions to scale solutions effectively.

Artificial intelligence (AI) has matured from promise to practice across select clinical tasks, yet real-world deployment continues to lag behind expectations¹. Responsible operationalization requires that ethical, technical, and organizational challenges be defined, tested, and governed in context, not merely described in principle. Feasibility studies help bridge this gap, but bedside evaluation still depends on how well systems are matched to local realities². In response, both the United States and the European Union have advanced regulatory pathways and collaborative initiatives that seek to make AI useable and safe at scale. Emerging strategies in low- and middle-income countries (LMICs) are moving in the same direction, but sector-specific guidance for health care remains limited and uneven^{3,4}. Although it is not the only factor that will affect the developmental trajectory of AI, local government capacity and infrastructure will have some influence on whether, and how, AI reaches patients. As such, it is insufficient to pay attention to model performance alone and, furthermore, local realities must closely inform government-led legislation and policies.

Why regional networks are necessary

As legislation and standards evolves, large collaborations have begun to coordinate trustworthy AI in health care. Networks such as the Trustworthy & Responsible AI Network (TRAIN) in the US⁵ and TRAIN-Europe⁶, as well as the United Nations-based AI for Good Global Summit⁷, and the African Union-based Digital Transformation Strategy for Africa⁸ illustrate how shared guardrails, post-deployment monitoring, and common language can accelerate safe adoption. However, the contexts in which these models will be used are not interchangeable. European policy emphasizes rights-based regulation and post-market surveillance, while the United States relies on a more decentralized, sector-specific approach. Low-resource regions must

navigate different constraints: heterogeneous or nascent regulation, fragmented data ecosystems, language under-representation, intermittent connectivity and power, and clinical workflows centered on primary care and public health rather than tertiary services. Transferring US or EU frameworks wholesale risks burdening already stretched systems with documentation and compliance demands that are misaligned with available capacity. We believe that regional networks built for low-resource settings can translate high-level principles into feasible steps, sequence requirements according to maturity, and prioritize use cases that reflect local disease burdens and service organization. Thus, for the purposes of this paper, we define a regional network as: a collaborative body, usually composed of locally initiated and/or geographically focused entities, that work together to address shared challenges and opportunities within a specified geographic area. In the context of AI, these networks involve data infrastructure, governance structures, and technical expertise to facilitate effective collaboration both within the region and with international partners.

Local health needs and contextual relevance

There are several applications of interest to low-resource settings, from diagnosis, epidemic surveillance, and telemonitoring, to screening and risk stratification. Operational success depends on whether models help with the problems that frontline clinicians and communities face every day. In many African countries, for example, health systems contend with both noncommunicable diseases and a spectrum of infectious threats, where diagnosis and screening applications could assist. Lassa fever in West Africa illustrates the diagnostic challenge created by nonspecific early symptoms that overlap with malaria or typhoid; a substantial proportion of infections may be mild or asymptomatic, complicating case finding and triage⁹.

¹Department of Adult Intensive Care, Erasmus MC, Rotterdam, The Netherlands. ²Erasmus MC Datahub, Erasmus MC University Medical Center, Rotterdam, The Netherlands. ³University of Toronto, Toronto, Ontario, Canada. ⁴Department of Medicine, Stellenbosch University, Stellenbosch, South Africa. ⁵Delft University of Technology, Delft, The Netherlands. ⁶Department of Biostatistics and Bioinformatics, Duke University School of Medicine, Durham, NC, USA.

✉ e-mail: m.vangenderen@erasmusmc.nl

Fig. 1 | Status of National AI Strategies on the African continent (2025). The map was created, using tools Canva and Flourish, by the authors based on publicly data available from the AI Policy Lab's overview of national AI strategies across Africa^{24,25}, with categories interpreted and classified by the authors according to national policy status. Colors represent national AI strategy status: legislated strategy, published strategy, strategy under development, commitment to develop a strategy, or no known strategy. The figure illustrates cross-country variation in national AI policy readiness.



Similarly, epidemic surveillance applications would be invaluable for containing the spread of another local virus known as Nipah virus. The virus causes severe and often fatal encephalitis. In parts of South and Southeast Asia, outbreaks of Nipah virus have shown high case-fatality rates in some settings, and early presentations again resemble common febrile illnesses¹⁰. Diagnostic support that ignores these epidemiologic realities, or that is trained on data from settings where these conditions are rare, will underperform at the bedside. Conversely, focused tools aligned to local priorities can succeed under real constraints. Moreover, up to 75% of cases result in fatality¹¹. Therefore, AI models that can adequately assist in diagnosing people with Nipah virus should be familiarized with Nipah virus prevalence and data on its diagnosis.

Feasibility work on AI-assisted diabetic retinopathy screening in Rwanda suggests that targeted diagnostics embedded in care pathways can be acceptable and workable¹². AI-enabled telemedicine programs in India demonstrate how standardized intake and triage can extend reach when integrated with national digital health infrastructure¹³. We believe these cases demonstrate the direction in which AI must move in order to remain contextually relevant and genuinely responsive to the health burdens across the globe. A one-size-fits-all approach will not do.

Given the burden of disease in many LMICs, we believe AI could additionally play a crucial role in pandemic preparedness and response, epidemiological tracking, and disaster management. AI could, for instance, be leveraged for early warning systems and public health monitoring to minimise the impact of global health disasters disproportionately affecting low-resource settings. Recent responsible AI operationalization has been exemplified by initiatives such as Rwanda's establishment of Africa's first AI Scaling Hub—a government-backed platform designed to identify, refine, and scale AI solutions across health, agriculture, and education, including AI-powered telemedicine, supply-chain forecasting, and diagnostic tools tailored for underserved communities¹⁴. It is precisely initiatives such as these that could act as a blueprint for other countries in the region.

Legislation, infrastructure, and the path from principles to practice

The regulatory landscape for AI in healthcare is evolving rapidly, with regions like the EU and the US taking significant steps toward establishing comprehensive frameworks for transparency, risk management, and system monitoring. However, low-resource settings face unique challenges in developing context-specific AI governance frameworks. While international frameworks may serve as useful benchmarks, locally-led approaches are needed to reflect national priorities, infrastructural realities, and institutional capacities. In this commentary, we focus primarily on Africa, given the sheer size and political complexity of the continent, where diverse regulatory environments and limited resources highlight the challenges of AI regulation in healthcare.

The EU AI Act and emerging US policy actions have clarified expectations for transparency, risk management, and post-deployment oversight¹⁵. Many countries in Africa and other low-resource regions are developing strategies and guidance, yet comprehensive, binding health-AI regulation is still uncommon and capacity for enforcement varies widely (Fig. 1). This regulatory heterogeneity interacts with infrastructural constraints. Despite the fact that several African countries have put in place the first steps toward establishing national AI strategies to guide development¹⁶, these alone are insufficient. Electronic health record coverage may be limited, standards unevenly applied, and privacy and security baselines difficult to maintain. Connectivity and power can be intermittent, and compute resources are often concentrated in capitals or referral hospitals rather than distributed to primary-care facilities. Regional networks can convert broad principles into executable playbooks that account for these realities. This includes establishing light-weight registries, fit-for-purpose oversight processes, and post-deployment monitoring that is credible and auditable without being prohibitive. It also means aligning deployment with connectivity tiers and pursuing an “offline-first” posture where needed, while planning for progressive enhancement as capacities improve.

Table 1 | Challenges for the safe and responsible adoption of AI in low-resource settings: a practical approach

Category	Question	Purpose	Rationale	Solution
Governance	Are actionable regulations or strategies in place for safe, responsible health-AI adoption?	Provide accountable oversight with continuous monitoring and evaluation.	Governance frameworks are critical for establishing standards to guide AI implementation, ensuring accountability of all stakeholders.	Establish legal frameworks, an AI registry, clear accountability, incident reporting, and post-deployment review cycles.
Data protection and data bias	Are there regulatory and protective structures in place to safeguard medical data and prevent bias in AI-models?	- Ensure the ethical collection, storage and use of medical data and mitigate cybersecurity risks. - Create participation and data representation in global AI research and model development.	Data security and risk of bias are major barriers in low-resource settings, where data laws may be insufficient, and bias in models could lead to inequitable outcomes.	- Establish strong, transparent data privacy regulations. - Put in place cybersecurity measures in organizations that handle medical data (e.g. private clinics, hospitals, mobile labs). - Encourage rigorous testing, ongoing monitoring and evaluation of model performance. - Introduce data infrastructure for the collection and standardization of clinical data.
Local/regional contexts	Are local health needs and priority AI use cases clearly defined?	Improve awareness about the specific local needs of populations and possible AI solutions.	Local awareness of disease burdens is crucial for ensuring that AI models are relevant and democratically target the needs of the population.	Foster collaborations between the technology and frontline healthcare sectors through establishing pilot projects and research networks.
Community-driven AI for primary healthcare	Are AI tools being targeted towards primary care, or do models focus insufficiently on these challenges?	Given that low-resources settings struggle at a grass roots level to provide primary care, AI models ought to target these health burdens first and foremost.	Primary care is often underfunded or of unequally distributed. AI tools must thus either directly contribute to primary care or at least not come at a cost to primary care.	Support community-health-worker programs; prioritize screening, adherence support, and self-testing guidance integrated with referral pathways.
Indigenous knowledge and traditional medicine	Are there existing traditional healing practices that must be considered by AI systems in order to be effective?	Recognizing that traditional healing practices will affect predictive modeling allows AI models to be responsive to these idiosyncrasies.	Traditional healing is often a big part of many cultures. AI implementation must therefore be sensitive and inclusive so ensure trust and acceptance.	AI developers need to assess the potential for traditional healing practices to affect their predictive modelling in order to account for this in their models.
Investment and affordability	Are there enough local and national investments to make the integration of AI affordable?	Make the integration of AI affordable, scalable, and equitable, while making local innovations sustainable.	Financial barriers are arguably the major barrier in low-resource settings. Thus, AI adoption requires sustainable financial solutions, particularly in these regions.	- Acquire funds through public-private partnerships (PPP). - Support the creation of open science and data sharing to prevent monopolies on clinical AI.
Infrastructure, energy, and data	Can sites run AI reliably given power, compute, and connectivity constraints?	Enable the technological adoption of AI by enlarging the capacity for data storage, internet connection and computing power.	Unlike in high-resource settings, network, energy, and connectivity can prove to be major barriers to AI solutions.	Plan for connectivity tiers and offline-first operation; deploy edge compute where needed; invest in resilient power and basic data hubs.
Public-private partnerships (PPPs)	- Is there the possibility for successful collaborations between governments, NGOs and tech companies? - Have the possible collaborations been exhaustively explored, taking into account costs and risks? - Furthermore, have these risks been weighed up against costs and benefits of domestic financing?	PPPs can bridge both financial and technical gaps that might exist in particularly limiting ways.	While PPPs can be one source of potential funding, they must be very carefully evaluated for risks and costs and, additionally, be fully weighed against other potential solutions.	- Provide exemplars of equitable, successful collaborations that have resulted in health gains for burdened populations, and encourage further collaborations of this nature. - Create networks and consortia that facilitate these collaborations.
Educational needs	Are healthcare professionals as well as policy makers equipped with the relevant AI literacy to ensure responsible implementation and good governance?	Upskilling of key actors responsible for implementation, narrowing the expertise divide that exists between low-resource and high-resource settings.	Missing AI knowledge can result in a lack of trust as well as an inability to leverage AI solutions. Thus, educational needs must be developed alongside AI solutions.	Implement, and monitor, education and training programs to enhance local expertise.
Trust and expertise	Are healthcare professionals capable of handling AI software and do they trust it?	Enhance AI adoption by exploring productivity improvements, training, support, and workflow integration, and fostering a culture of innovation.	Trust goes hand-in-hand with adoption. Thus, building confidence in AI systems will enhance acceptance among healthcare workers and likely improve outcomes for patients.	Provide in-workflow support and feedback loops; monitor adoption and outcomes; iterate based on user-reported issues.

Category: The key thematic area addressing a specific challenge related to the adoption and operationalization of AI in low-resource settings.

Question: The relevant question that highlights the main issue to be addressed within the category.

Purpose: The desired outcome of addressing the challenge in each category—i.e., what needs to be achieved to ensure responsible AI adoption.

Rationale: A brief explanation of why this category is critical for AI adoption in low-resource settings, outlining the underlying issues and context.

Solution: Practical actions that can be taken to address the challenge, focusing on concrete, actionable steps.

Moreover, it is unfortunately not sufficient to simply import legislative guidance from the EU or US. Both frameworks may offer important insights, but they also present challenges in their application to low-resource settings. In particular, the EU's stringent data protection and accountability requirements may pose barriers in regions with limited infrastructure, while the US's more market-driven and flexible regulatory approach, which places greater emphasis on enabling private-sector development, could face challenges in ensuring equitable outcomes for more vulnerable populations. The broad comparison of these frameworks is therefore an indication that developing a regionally-specific approach, while necessary, is no small task.

Data inclusion as a prerequisite for equity and performance

Model fairness and generalizability are functions of who and what are represented in training data. Clinical text corpora used for language-model development still include little or no data from many African countries. Major world languages dominate available datasets while large portions of linguistic diversity remain under-represented¹⁷. In clinical imaging and dermatology, darker skin tones are frequently a minority in training sets, despite well-documented disparities in detection and outcomes that are driven by both biology and access¹⁸. The result is predictable: models trained on narrow slices of the world struggle when deployed elsewhere. Regional networks provide a mechanism for moving from critique to construction by brokering data partnerships, setting inclusion targets that can be measured, and prioritizing privacy-preserving methods that enable participation without sacrificing safety. A good example of this is a recent attempt at using federated learning for tuberculosis diagnosis in low-resource settings across Africa, enabling collaborative model development while maintaining patient data privacy¹⁹. The aim is not to replicate the largest datasets, but to build representative, high-value corpora and benchmark tasks that capture local disease patterns, languages, and workflows. This is especially important in primary care and public health, where simple variables and structured processes can deliver outsized value if consistently captured and shared.

From pilots to sustained value

Experience shows that pilots often gain initial traction despite their frequent inability to scale successfully²⁰. This challenge, known as “pilotitis”²¹, is often compounded by donor-funded, short-term initiatives that prioritize pilot testing without sufficient planning for long-term sustainability. To progress beyond one-off demonstrations, deployment must be coupled to maintenance, monitoring, and clear measures of value. For frontline clinicians, this includes whether a tool reduces time, enables safer triage, or improves referrals and follow-up under real staffing constraints. For patients, it includes acceptability, privacy protection, and the perceived legitimacy of the system. For ministries and payers, it includes coverage, continuity, and the ability to detect and correct performance drift over time. Regional networks can make these expectations explicit and comparable by defining minimum reporting elements, shared indicators, and simple pathways for incident disclosure and remediation. When combined with open standards and public-interest clauses in partnerships, these practices reduce the risk of vendor lock-in and increase the likelihood that systems can be maintained locally.

Towards a regional approach

As a global technology with cross-border effects, AI demands better representation in research, data inclusion, and model development. Yet, reports like the latest AI Index Report from the Institute for Human-Centered AI, which aim to highlight global AI trends, fail to adequately include African countries or African authors in their report²¹. Similarly, the Artificial Intelligence Index Report 2024²² – presented as a comprehensive global assessment, lacks representation from low-resource settings among its staff, researchers, and steering committee. On the other hand, we acknowledge the important work contained in the Artificial Intelligence and Democratic Values 2025 report²³, wherein low-resource settings are included. However, this reporting standard is not indicative of the global trend,

where lack of inclusion underscores the need for regionally-specific approaches that address local challenges and produce tailored solutions (Table 1).

The challenge and need for AI operationalization have led to the establishment of international consensus guidelines, checklists, and several networks geared towards responsible AI for healthcare. However, for AI to be operationalized responsibly on a global scale, it must include the voices and needs of all regions, particularly those in low-resource settings. This requires not only the development of local AI regulations and data infrastructures but also the active inclusion and participation of local healthcare professionals in AI development. It should be noted, therefore, that the solutions presented in Table 1 are not value-neutral. They are underpinned by ethical principles specifically important to low-resource settings, such as equity, inclusivity, and long-term sustainability. Local solutions, however, should ideally align with existing standards global standards for AI in order to prevent isolated solutions, which might make collaborations difficult. Thus, where local solutions diverge from global standards, these must be well justified from an ethical perspective. Only through such efforts can AI be truly operationalized in a way that is inclusive, ethical, and beneficial to all.

Conclusion

Responsible AI becomes real only when patients, clinicians, and health systems can use it safely and reliably where care is delivered. Low-resource settings deserve solutions designed for their realities, not adapted as an afterthought. Regional implementation networks provide the connective tissue required to translate principles into practice: credible oversight that matches capacity, data inclusion that reflects populations and languages, workforce development that builds stewardship, and infrastructure choices that work when connectivity and power do not. Supporting these networks through open standards, transparent partnerships, and sustained investment is both an ethical obligation and a practical route to models that generalize, perform, and earn trust.

Data availability

No datasets were generated or analysed during the current study.

Received: 2 December 2025; Accepted: 27 August 2026;

Published online: 21 September 2026

References

1. Berkhout, W. E. M. et al. Operationalization of artificial intelligence applications in the intensive care unit: a systematic review. *JAMA Netw. Open* **8**, e2522866 (2025).
2. van de Sande, D. et al. Operationalization of artificial intelligence to assist in surgical discharge: a feasibility case study. *NEJM AI* **3**, Alcs2401132 (2025).
3. Suleiman, S., Olén, D. & Bezuidenhout, M. AI regulation and policy in Africa. <https://www.dentons.com/en/insights/articles/2024/june/13/ai-regulation-and-policy-in-africa> (2024).
4. Xu, J., Lee, T. & Goggin, G. AI governance in Asia: policies, praxis and approaches. *Commun. Res. Pract.* **10**, 275–287 (2024).
5. Embí, P. J., Rhew, D. C., Peterson, E. D. & Pencina, M. J. Launching the Trustworthy and Responsible AI Network (TRAIN): a consortium to facilitate safe and effective AI adoption. *JAMA* **333**, 1481–1482 (2025).
6. van Genderen, M. E., Kant, I. M. J., Tacchetti, C. & Jovinge, S. Moving toward implementation of responsible artificial intelligence in health care: the European TRAIN initiative. *JAMA* **333**, 1483–1484 (2025).
7. ITU. *AI for Good Global Summit*, 2026 (ITU, 2025).
8. African Union. *The Digital Transformation Strategy for Africa (2020–2030)* (African Union, 2020).
9. WHO. *Lassa Fever* (WHO, 2025).
10. Ang, B. S. P., Lim, T. C. C. & Wang, L. Nipah virus infection. *J. Clin. Microbiol.* **56**, 01875–01817 <https://doi.org/10.1128/jcm> (2018).

11. Epstein, J. H., Field, H. E., Luby, S., Pulliam, J. R. C. & Daszak, P. Nipah virus: impact, origins, and causes of emergence. *Curr. Infect. Dis. Rep.* **8**, 59–65 (2006).
12. Whitestone, N. et al. Feasibility and acceptance of artificial intelligence-based diabetic retinopathy screening in Rwanda. *Br. J. Ophthalmol.* **108**, 840–845 (2024).
13. Nwankwo, E. I., Emeihe, E. V., Ajegbile, M. D., Olaboye, J. A. & Maha, C. C. Integrating telemedicine and AI to improve healthcare access in rural settings. *Int. J. Life Sci. Res. Arch.* **7**, 59–77 (2024).
14. Gacuruzwa, E. *Rwanda Launches \$17.5 Million AI Scaling Hub to Lead Africa's Digital Future* (Further Africa, 2025).
15. Busch, F. et al. AI regulation in healthcare around the world: what is the status quo? *medRxiv*, <https://www.medrxiv.org/content/10.1101/2025.01.25.25321061v1> (2025).
16. Akbaraliabad, H. et al. *The Utility of Generative AI in Advancing Global Health*, vol. 2 Alp2400875 (Massachusetts Medical Society, 2025).
17. Joshi, P., Santy, S., Budhiraja, A., Bali, K. & Choudhury, M. The state and fate of linguistic diversity and inclusion in the NLP world. In *Proc. Annu. Meet. Assoc. Comput. Linguist.* 6282–6293 (2020).
18. Daneshjou, R. et al. Disparities in dermatology AI performance on a diverse, curated clinical image set. *Sci. Adv.* **8**, eabq6147 (2022).
19. Fabila, J., Garrucho, L., Campello, V. M., Martín-Isla, C. & Lekadir, K. Federated learning in low-resource settings: a chest imaging study in Africa--Challenges and lessons learned. Preprint at <https://arxiv.org/abs/2505.14217> (2025).
20. Greenhalgh, T. et al. Beyond adoption: a new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. *J. Med. Internet Res.* **19**, e8775 (2017).
21. Bhatia, A., Matthan, R., Khanna, T. & Balsari, S. Regulatory sandboxes: a cure for mHealth pilotitis? *J. Med. Internet Res.* **22**, e21276 (2020).
22. The Stanford Institute for Human-Centered AI. *Artificial Intelligence Index Report 2025* (The Stanford Institute for Human-Centered AI, 2025).
23. Maslej, N. et al. Artificial Intelligence Index. AI Index Report 2024 (Stanford Institute for Human-Centered Artificial Intelligence, 2024); stanford.edu.
24. Center for AI and Digital Policy. *Artificial Intelligence and Democratic Values Index 2025* (Center for AI and Digital Policy, 2025).
25. AI Policy Lab. *National AI Strategies in Africa*, vol. 2026 (AI Policy Lab, 2025).

Acknowledgements

This research is part of the Responsible and Ethical AI for Healthcare Lab (REAIHL). No funding was received for this research.

Author contributions

E.J.S., M.V.G., J.V.D.H., and D.G. conceptualized the idea and arguments for this paper. E.J.S. and J.H. wrote the first version. J.S. and K.M. edited and provided extensive commentary. M.V.G. and E.J.S. wrote the final version.

Competing interests

D.G. reported receiving grants from SAS Software Inc. outside the submitted work. No other disclosures were reported. All other authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Michel E. van Genderen.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026