Equity by design in artificial intelligence for primary care in low-income and middle-income countries

Lancet Prim Care 2026

Published Online
https://doi.org/10.1016/j.lanprc.2026.100135

Primary health care is the initial point of contact for most people seeking care. Clinicians need to make rapid decisions on undifferentiated symptoms while delivering prevention and chronic care, often with scarce diagnostics, little supervision, and high cognitive load. In many low-income and middle-income countries (LMICs), three constraints reinforce one another: scarcity of workforce and resources, quality gaps, and data systems that rarely return value to the point of care.

Workforce scarcity is structural: a global shortfall of roughly 10 million health workers is projected by 2030, concentrated in regions with the highest disease burdens.[1] Quality gaps compound scarcity; the *Lancet Global Health* Commission estimated that 8·6 million deaths per year in 137 LMICs are associated with inadequate access to quality care, mostly among people who seek care but receive poor-quality services.[2] In primary health care, these issues in access and quality appear as missed diagnoses, inappropriate prescribing, and insufficient follow-up driven by workload and missing information.

Data is the third bottleneck. Many primary health-care settings still rely on paper or fragmented digital tools, creating administrative load. Across five LMICs, facilities were required to maintain numerous registers and recurring reporting forms.[3] When digitisation does not create feedback loops, it becomes clerical modernisation. Fragmentation also produces health-data poverty—the under-representation of languages, populations, and care settings in high-quality digitised data.[4] Health-data poverty is an equity-failure mode for artificial intelligence (AI). Models can be least safe where need is greatest, and performance claims rarely travel across differences in disease patterns, diagnostics, and referral capacity.

These constraints now collide with global health austerity. Deaths among children younger than 5 years were projected to rise in 2025 to approximately 4·8 million,[5] while official development assistance declines are hitting health systems hard. As a result, doing more with less should focus on making primary health care more reliable and equitable, not using tools that increase risk.

AI can serve as a reliable force multiplier, not a replacement (figure). AI's promise in primary care in LMICs is not to replace clinicians, but to increase the effective supply of safe clinical support. When used well, AI can act as an expert in the pocket, guiding structured assessment, danger signs, and guideline-aligned next steps tailored to available resources. It can also reduce administrative burden by drafting notes and summarising existing data, and it can make supervision more scalable by helping clinicians support and mentor frontline workers.

Consumer self-care and care navigation can also serve as a digital initial point of contact. In congested systems, households often pay in time and money to discover that home care or a nearby community service is sufficient. When done well, navigation tools can reduce unnecessary facility demand and identify danger signs earlier. But triage and advice are crucial for safety: intended use, escalation thresholds, and local-language performance need to be explicit and tested.

Compared with earlier rule-based clinical decision support, which delivered modest average gains and often added workflow friction,[6] large language models offer a more flexible interface for natural-language inputs and multitask support.[7] For primary health care, equity will depend on hybrid systems—large language model

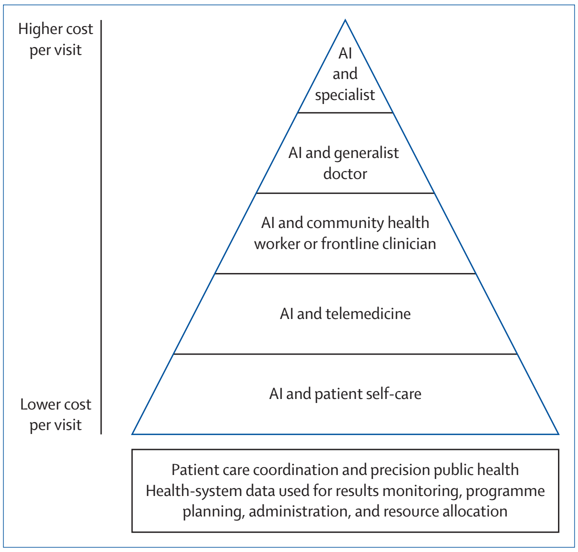

***Figure:* AI-enabled primary health care**
AI has the potential to improve expertise at each level of the health-care system. Over time, more interaction can be handled at lower levels for lower cost per interaction or per patient, increasing efficiency and access. AI=artificial intelligence.

interfaces coupled with structured, guideline-derived logic and locally validated workflows—so that usability improves without losing accountability.

Moving from hype to impact requires institutionalised reliability. Equitable AI for primary care in LMICs should pass five tests. These tests are a practical checklist for ministries, implementers, and funders. They shift attention from headline model performance to what matters in primary health care: safety under routine conditions, usability in real workflows, inclusion of underserved groups, and the ability to detect and correct harm over time. Failure on any one test should slow down scale-up; passing them usually implies investment in data, integration, and monitoring after deployment, not just software procurement.

First, does the tool relieve the clinician's hardest moment? A rural nurse might treat more than 70 patients daily, facing high-risk decisions without sufficient diagnostics[8] or complete records. Equity-driven AI should guide danger-sign assessments and actionable, guideline-based decisions. Prediction tools that do not support action only add to workload; solutions need to address actual constraints, not just showcase technology.

Second, does the system make invisible patients visible? Fragmented records make patients who move, miss follow-up, or use multiple entry points clinically invisible. AI trained on partial data will reproduce these gaps. Equity requires interoperable identifiers and minimum datasets that work across paper and digital flows, support local languages, and return insights to the point of care, especially for rural, mobile, and marginalised populations.

Third, does the evidence reflect real clinical conditions? Benchmark accuracy rarely captures the realities of primary health care: connectivity gaps, little supervision, constrained diagnostics, and uneven referral pathways. AI might improve scores yet increase consultation time or inappropriate referrals. Equity requires pragmatic evaluation under routine conditions, with outcomes, workflow measures (eg, time or number of referrals), and disaggregated safety by geography, language, and gender, not accuracy alone. Reporting standards, such as Consolidated Standards of Reporting Trials (CONSORT)-AI, should be the minimum.[9] What is key to understand is not whether AI knows more than a human, but rather whether the form of augmented knowledge is usable to improve care effectively at the point of service.

Fourth, can decisions be understood, reviewed, and corrected? Frontline workers remain accountable. Equity depends on verifiability: clear intended use, documented performance, traceable outputs, and audit trails that enable review, appeal, and correction. Systems should communicate uncertainty and allow override. Transparency and governance matter more than sophistication, and WHO guidance provides a baseline.[10]

Fifth, does the tool protect the workforce, or expose it? Primary care workers already absorb the consequences of system failures. Introducing AI without training, supervision, and defined accountability can shift liability onto clinicians already under strain and increase moral distress when tools are wrong but unavoidable. Equitable implementation requires investment in clinical and digital literacy, supervision structures, and safety roles to monitor errors, bias, and large language model output drift. Without a workforce strategy, there is no safety strategy.

AI could enable a step change in primary health care in LMICs, but impact will depend less on model novelty than on data foundations, governance, workforce capacity, and evaluation for safe implementation. The goal is a more reliable and humane primary health care encounter, wherein clinicians spend more time on patients, not paperwork, and households have safer guidance closer to home. Equity is the design specification for scalable impact. Equity by design means safety, inclusion, and accountability are specified before scale, not retrofitted after harm.

ZB and MM are employees of the Gates Foundation. SG is the CEO of Qhala. AB is the Director of Ariadne Labs. During the preparation of this work, the authors used ChatGPT (version 5.2) to edit the language and grammar and to identify references. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

*Zameer Brey, Shikoh Gitau, Melissa Miles, *Asaf Bitton*
abitton@ariadnelabs.org

Gates Foundation, Seattle, WA, USA (ZB, MM); Qhala, Nairobi, Kenya (SG); Ariadne Labs, Harvard T H Chan School of Public Health and Brigham and Women's Hospital, Boston, MA 02145, USA (AB)

1 Boniol M, Kunjumen T, Nair TS, Siyam A, Campbell J, Diallo K. The global health workforce stock and distribution in 2020 and 2030: a threat to equity and "universal" health coverage? *BMJ Glob Health* 2022; **7:** e009316.

2 Kruk ME, Gage AD, Arsenault C, et al. High-quality health systems in the Sustainable Development Goals era: time for a revolution. *Lancet Glob Health* 2018; **6:** e1196–252.

3 Siyam A, Ir P, York D, et al. The burden of recording and reporting health data in primary health care facilities in five low- and lower-middle income countries. *BMC Health Serv Res* 2021; **21** (suppl 1)**:** 691.

4 Paik KE, Hicklen R, Kaggwa F, et al. Digital Determinants of Health: health data poverty amplifies existing health disparities—a scoping review. *PLoS Digit Health* 2023; **2:** e0000313.

5 Gates Foundation. 2025 Goalkeepers report. We can't stop at almost. https://www.gatesfoundation.org/goalkeepers/report/2025-report/ (accessed Feb 20, 2026).

6 Kwan JL, Lo L, Ferguson J, et al. Computerised clinical decision support systems and absolute improvements in care: meta-analysis of controlled clinical trials. *BMJ* 2020; **370:** m3216.

7 Singhal K, Azizi S, Tu T, et al. Large language models encode clinical knowledge. *Nature* 2023; **620:** 172–80.

8 Yadav H, Shah D, Sayed S, Horton S, Schroeder LF. Availability of essential diagnostics in ten low-income and middle-income countries: results from national health facility surveys. *Lancet Glob Health* 2021; **9:** e1553–60.

9 Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *Lancet Digit Health* 2020; **2:** e537–48.

10 WHO. Regulatory considerations on artificial intelligence for health. Oct 19, 2023. https://www.who.int/publications/i/item/9789240078871 (accessed Feb 20, 2026).