

The AI Health Arms Race: A Critical Perspective on Big Tech and the Widening Global Health Equity Gap

Mohamed Mustaf Ahmed ¹, Zhinya Kawa Othman ²

¹Faculty of Medicine and Health Sciences, SIMAD University, Mogadishu, Somalia; ²Faculty of Pharmaceutical Sciences, Chulalongkorn University, Bangkok, Thailand

Correspondence: Mohamed Mustaf Ahmed, Email momustafahmed@outlook.com

Abstract: The first quarter of 2026 witnessed an unprecedented convergence, with OpenAI, Anthropic, Microsoft, Google, and Apple launching or advancing dedicated artificial intelligence health platforms. ChatGPT Health, Claude for Healthcare, Copilot Health, Med-Gemini, and Apple Health+ collectively represent a paradigm shift toward AI-mediated personal health management, integrating electronic health records, wearable device data, and conversational AI in privacy-isolated environments. However, these tools are primarily designed for high-income country markets, with limited infrastructure, insufficient multilingual support beyond dominant global languages, and minimal cultural adaptation for low- and middle-income countries. This commentary critically examines the emerging AI health chatbot landscape through the lens of global health equity, analyzing structural barriers, including data poverty, regulatory vacuums, and the risks of data colonialism, whereby large technology corporations extract health data from populations in low- and middle-income countries without proportionate benefit sharing or local capacity building. We propose policy recommendations spanning international governance, national regulatory development, mandatory multilingual content, pre-market clinical safety evaluations, and multilateral financing of digital health infrastructure. We further discuss the strategic responsibilities of both high-income country technology corporations and governments in low- and middle-income countries in bridging this divide. Without deliberate equity-centered governance, the AI health arms race risks widening, rather than narrowing, the global health divide.

Keywords: artificial intelligence, health equity, large language models, digital health, low- and middle-income countries

Introduction

The first quarter of 2026 witnessed an unprecedented convergence of major technology companies launching dedicated artificial intelligence health platforms aimed at consumers. In January 2026, OpenAI introduced ChatGPT Health, a dedicated space within ChatGPT that allows users to connect medical records and wellness applications to receive personalised health responses.¹ Within days, Anthropic launched Claude for Healthcare, offering Health Insurance Portability and Accountability Act (HIPAA)-ready tools for providers, payers, and patients, with connectors to medical databases and coding systems.² In March 2026, Microsoft debuted Copilot Health, integrating wearable data from over 50 devices and electronic health records (EHR) from more than 50,000 United States hospitals.³ Google has expanded its Med-Gemini family of models for clinical applications and partnered with b.well Connected Health for health data connectivity, while Apple continues to develop AI-powered health features for integration into its Health application ecosystem.⁴

The term “arms race” in this context draws on its established meaning in political science and international relations, describing a competitive dynamic in which multiple actors escalate their investments and capabilities in pursuit of a strategic advantage, often resulting in outcomes that undermine collective welfare. In the domain of AI-driven health platforms, this competitive escalation among technology corporations risks prioritizing market dominance and proprietary data ecosystems over equitable access, creating a landscape in which innovation is concentrated in high-income

markets while low- and middle-income countries (LMICs) are left further behind.⁵ This framing highlights the urgency of governing AI health technologies to prevent the entrenchment of global health inequities.

This rapid commercialization reflects the enormous consumer demand for AI-mediated health information; OpenAI reported that over 230 million people ask health and wellness questions on ChatGPT each week, and Microsoft handles more than 50 million health queries daily across its consumer products.^{1,3} However, the simultaneous launch of these platforms raises critical questions about global health equity, data governance, clinical safety, and the potential for AI health tools designed in high-income settings to deepen existing disparities in LMICs.

This commentary adopts a public health equity and ethical governance lens, integrating perspectives from global health policy, technology governance, and distributive justice to critically evaluate how the competitive deployment of AI health platforms may exacerbate or ameliorate existing health disparities.⁶ By situating the analysis within this framework, we aim to move beyond descriptive accounts of technological innovation toward a critical examination of the structural, regulatory, and ethical dimensions that determine whether these technologies serve global health or deepen global inequity.

This commentary examines the emerging AI health chatbot landscape through the lens of global health equity, identifies the structural barriers that may prevent equitable benefit sharing, and proposes policy-oriented recommendations for ensuring that these technologies serve populations in need of improved health access. The structural barriers examined include limited digital infrastructure, insufficient health data systems, regulatory gaps in data protection and AI governance, language and cultural barriers in AI content delivery, and the extractive dynamics of data collection by major technology corporations operating across borders, whereby health data flows from LMICs to high-income country platforms without commensurate investment in local health systems.⁷

The Emerging Landscape of AI-Powered Consumer Health Platforms

These platforms share common architectural features, including EHR integration, wearable device connectivity, privacy-isolated health spaces, and commitments that health data will not be used to train foundation models (Table 1). ChatGPT Health was developed with more than 260 physicians across 60 countries, connecting to Apple Health, MyFitnessPal, and Function.¹ Claude for Healthcare provides enterprise-grade connectors to the CMS Coverage Database, ICD-10 codes, and PubMed.² Copilot Health aggregates data from 50+ wearable devices and hospitals through HealthEx, with answers verified by more than 230 physicians from 24 countries.³ Google's Med-Gemini demonstrates capabilities in radiology, pathology, and genomic risk prediction.⁸ Reports also claim that Apple has been developing Health+ with AI coaching and nutrition tracking, although it has scaled back toward incremental releases.⁹ The competitive intensity is captured by Microsoft's characterisation of this work as steps toward "medical superintelligence".¹⁰

The scale of investment and speed of deployment underscore the arms race dynamic. These corporations are not merely developing health tools in parallel; they are actively competing for dominance in health data ecosystems, clinician

Table 1 Overview of Major AI Health Platforms Launched in 2025–2026

Platform	Developer	Key Health Features	Source
ChatGPT Health	OpenAI	EHR integration via b.well, Apple Health and MyFitnessPal connectivity, dedicated health conversation space, physician-collaborated development	[1]
Claude for Healthcare	Anthropic	HIPAA-ready enterprise tools, CMS and ICD-10 connectors, PubMed integration, HealthEx EHR access, FHIR development skills	[2]
Copilot Health	Microsoft	Data from 50+ wearable devices, EHRs from 50,000+ US hospitals via HealthEx, Harvard Health answer cards, provider directory search	[3]
Med-Gemini and Health AI	Google	Multimodal medical models for radiology, pathology, and genomics; Personal Health LLM for wearable sensor data interpretation; b.well partnership	[8]
Health app AI features and planned Health+	Apple	AI health coaching agent, nutrition and meal tracking, integration with Apple Watch biometric sensors, planned Siri health intelligence in iOS 27	[9]

partnerships, and consumer trust. This competitive escalation carries measurable consequences for equity, as the design priorities, language support, and regulatory compliance of these platforms are overwhelmingly calibrated to high-income, English-speaking markets.⁵ The concentration of AI health innovation within a small number of technology corporations raises concerns about market power, data monopolization, and the capacity of governments, particularly in LMICs, to regulate these actors effectively.

Structural Barriers to Equitable Global Access

Despite their potential, these platforms reflect the priorities of high-income nations. EHR integration is almost exclusively limited to US hospitals, and wearable ecosystems presuppose consumer purchasing power, which is unavailable in LMICs. The World Health Organization (WHO) has emphasized that LMICs currently use only approximately 5% of their available health data, with governance, infrastructure, and security remaining significant barriers.¹¹ Scholars have introduced the “AI deployment paradox,” wherein the conditions that necessitate AI interventions, including data poverty and structural inequities, simultaneously undermine their effectiveness.¹² The digital divide encompasses not only infrastructure gaps but also disparities in digital literacy, language support, and cultural appropriateness of health information delivered through these platforms.

These companies collect vast quantities of health data through multiple channels, including wearable devices, EHR integrations, conversational interactions, and third-party application programming interfaces. In many LMICs, this data extraction occurs in the absence of comprehensive data protection legislation, allowing corporations to aggregate and utilize health data from populations that receive limited benefits from the resulting AI products. The asymmetry between data extraction and service provision represents a core dimension of the equity gap, as populations contributing data to improve AI models may lack access to the very platforms those models’ power.¹³

The integration of AI into primary health care in LMICs faces additional barriers including data privacy concerns, unresolved data ownership questions, and economic constraints.⁷ The global AI divide risks reinforcing disparities across healthcare, education, and governance, and neglecting low-income countries in AI discussions contradicts the principles of distributive justice.¹⁴ The current arms race risks creating a new form of technological colonisation, in which data flows from LMICs to high-income country corporations without equitable benefit-sharing.¹²

Empirical evidence supports these concerns. A study assessing the generalizability of AI clinical models across hospitals in the United Kingdom and Vietnam found significant performance degradation when models trained on high-income country data were deployed in LMIC settings, highlighting the risks of assuming transferability without local validation.¹⁵ Furthermore, researchers have documented how digital health platforms operating in the Global South may embed Western medical epistemologies and professional identities, marginalising local clinical practices and knowledge systems.¹³ These findings suggest that the current trajectory of AI health deployment may not only fail to close the global health equity gap but may actively widen it by imposing externally developed standards without adequate contextualization of the local context.

Clinical Safety, Data Sovereignty, and the Regulatory Vacuum

A fundamental concern is the risk of hallucinations, in which large language models generate information that appears coherent but is factually incorrect. A systematic review has documented that large language models (LLMs) risk producing plausible yet incorrect medical statements, requiring robust human oversight.¹⁶ Medical hallucinations include fabricated patient information and unsupported treatment recommendations, paralleling cognitive biases in clinicians but occurring without clinical accountability structures.¹⁷ This risk is compounded in LMIC settings, where health literacy may be lower and access to professionals for verification is limited, as summarized in [Table 2](#).

Consumer health AI applications fall outside HIPAA, meaning that users have no specific protections in the event of a data breach.²⁰ In LMICs, the regulatory vacuum is even more pronounced, with many countries lacking comprehensive data protection legislation or health AI governance frameworks. The WHO issued guidance in 2024 with more than 40 recommendations on large multi-modal models for health, yet the pace of commercial deployment far exceeds regulatory development.²¹ The aggregation of health records, wearable data, and conversation histories raises concerns about data colonialism, whereby high-income countries collect health data from countries lacking parallel protections.¹¹ The WHO

Table 2 Key Policy Challenges and Recommended Actions for Equitable AI Health Deployment

Policy Domain	Key Challenge	Recommended Action	Source
Regulatory governance	Consumer health AI falls outside HIPAA and most LMIC regulatory frameworks	Develop AI-specific health data protection legislation that mandates transparency, safety testing, and accountability	[18]
Data sovereignty	Risk of data colonialism as health data flows from LMICs to high-income country corporations	Establish national and regional data sovereignty frameworks governing cross-border health data transfers	[11]
Clinical safety	LLM hallucinations generate plausible but incorrect medical information without accountability	Require pre-market clinical safety evaluations with mandatory adverse event and hallucination rate reporting	[16]
Infrastructure equity	EHR integrations and wearable ecosystems limited to high-income markets	Incentivise open-standard interoperability and invest in LMIC digital health infrastructure through multilateral financing	[7]
Linguistic and cultural inclusion	AI health platforms predominantly support English and reflect Western clinical norms	Mandate multilingual support and culturally adapted content as conditions for LMIC market authorisation	[19]

ethics guidance has identified this risk, noting that commercial expansion may conflict with the interests of populations whose data contributes to platform improvement without proportionate access.¹⁸

Toward Equitable Governance and Global Impact

The WHO's Global Initiative on AI for Health has been working to harmonise governance standards, with particular attention to LMICs, advancing ethical, regulatory, and operational dimensions of health AI governance.¹⁹ The 2021 WHO report on Ethics and Governance of Artificial Intelligence for Health established the first international consensus on ethical norms, and the 2024 updated guidance expanded these principles for generative AI.¹⁸ Health digitalisation in LMICs requires coordinated investment in governance, infrastructure, and security.¹¹ The international community should establish a global framework for the pre-market assessment of consumer AI health tools, analogous to the International Medical Device Regulators Forum's guidance on good machine learning practice for software as a medical device.¹⁹

A recurrent theme in discussions of AI health equity is the need for transparency; however, transparency itself requires a precise definition to serve as an actionable policy objective. In the context of AI health platforms, transparency encompasses at least three distinct dimensions: algorithmic transparency, referring to the disclosure of how AI models process health data and generate recommendations; data transparency, concerning the openness about what health data is collected, how it is stored, and with whom it is shared; and institutional transparency, pertaining to the accountability structures and governance mechanisms that technology companies and regulatory bodies maintain.²² Each dimension entails distinct responsibilities for different actors in the supply chain. Governments bear the responsibility of establishing regulatory frameworks that mandate disclosures and accountability. Technology companies must ensure that their algorithms, data practices, and decision-making processes are subject to independent audits and public scrutiny. International organizations, including the WHO, should facilitate the development of harmonised transparency standards that bridge regulatory differences across countries.⁶ Without specifying which dimensions of transparency are being invoked and which actors bear responsibility, calls for transparency risk remaining aspirational rather than actionable in practice.

Governments in LMICs must adopt proactive strategies rather than remain passive recipients of externally designed technologies. Strategic priorities include investing in national digital health infrastructure, developing indigenous AI capacity through training programmes and research partnerships, establishing regulatory frameworks for AI health governance, and negotiating equitable terms for data sharing with international technology corporations.¹¹ Regional collaboration among LMICs can amplify individual country efforts by enabling shared regulatory standards, pooled training datasets that reflect local disease burdens and demographics, and collective bargaining power in negotiations with technology corporations.

International collaboration is equally essential in this regard. Multilateral organizations should facilitate technology transfer agreements, support the development of regional AI health centers of excellence, and ensure that global AI

governance discussions include meaningful representation from LMICs. The private sector, including technology corporations driving this competitive landscape, bears the responsibility of ensuring that their products are designed with global health equity as a core principle, not an afterthought. This includes committing to pre-deployment clinical validation in diverse populations and settings, investing in multilingual and culturally adapted content, and supporting local data governance frameworks that protect the interests of LMIC populations.⁵

The design paradigm must shift from extending high-income country products to LMIC markets toward co-designing platforms for diverse health challenges and cultural contexts. The scoping review literature on AI in the Global South has identified promising applications, including remote diagnostics, mobile health integration, and real-time monitoring in low-resource settings.²³ Open-source models, federated learning, and technology-LMIC partnerships could provide more equitable pathways, if governance frameworks ensure accountability. Governments in LMICs should ensure that AI health platforms operating within their jurisdictions meet minimum standards for local language support, cultural adaptation, and transparency regarding the limitations of AI-generated health information. Technology companies should conduct and publish pre-deployment clinical safety evaluations specific to the populations of the countries in which they operate rather than relying exclusively on benchmarks developed in high-income settings. Patients, community organizations, and civil society must participate in the design of technologies, develop new standards and demand transparency.

Conclusion

The simultaneous launch of AI health platforms by OpenAI, Anthropic, Microsoft, Google, and Apple represent a watershed moment in the healthcare sector. These platforms hold genuine promise for improving health literacy and patient-clinician communication. However, the current trajectory, characterized by US-centric integrations, premium wearable ecosystems, English-language dominance, and an undefined regulatory environment, threaten to widen these inequities. The AI deployment paradox, in which tools designed to reduce disparities may instead exacerbate them, is particularly salient. Technology corporations must recognize that their commercial strategies have public health consequences and that market expansion into LMICs carries an ethical obligation to ensure equitable access, clinical safety, and respect for local data sovereignty. Addressing these challenges requires international governance frameworks, national regulatory mechanisms, and industry commitments to co-design AI health tools with and for underserved populations. The WHO's Global Initiative on AI for Health provides a critical foundation; however, its recommendations must be translated into binding obligations supported by the financial and technical investments necessary for implementation. The key stakeholders in this effort include national governments, particularly health and technology ministries in LMICs; international organizations, such as the WHO and World Bank; technology corporations developing AI health platforms; civil society organizations advocating for health equity; and the research community, which must generate context-specific evidence to inform policy. This commentary contributes to the field by providing an integrated analysis that bridges technology governance, public health equity, and ethical considerations, offering a framework for stakeholders seeking to ensure that AI-driven health innovations serve all populations equitably. Without deliberate equity-centered governance, the AI health arms race risks producing a future in which the world's most sophisticated health intelligence tools are available primarily to those who need them the least.

Data Sharing Statement

No datasets were generated or analyzed in this study. All data referenced are publicly available through the cited sources.

Acknowledgments

The authors acknowledge the use of Paperpal AI <https://paperpal.com/> for its "Language Edit" and "Make Academic" features to improve clarity and readability. This assistance was limited to linguistic refinement; all analyses and interpretations are the sole responsibility of the authors.

Funding

This work received no specific funding from any funding agency in the public, commercial, or not-for-profit sectors.

Disclosure

The authors declare no conflicts of interest.

References

1. Introducing ChatGPT Health. OpenAI. Available from: <https://openai.com/index/introducing-chatgpt-health>. Accessed April 11, 2026.
2. Advancing Claude in healthcare and the life sciences. Anthropic. Available from: <https://www.anthropic.com/news/healthcare-life-sciences>. Accessed April 11, 2026.
3. Introducing Copilot Health. Microsoft. Available from: <https://microsoft.ai/news/introducing-copilot-health>. Accessed April 11, 2026.
4. Advancing cutting-edge AI capabilities. Google Health. Available from: <https://health.google/ai-models>. Accessed April 11, 2026.
5. Zaim R, Shaw JA. Toward responsible artificial intelligence in health: regulatory structures and power dynamics of the big tech industry in the United States. *Policy Studies*. 2024. doi:10.1080/01442872.2024.2400922
6. Berdahl CT, Baker L, Mann S, Nouri S. Strategies to improve the impact of artificial intelligence on health equity: scoping review. *JMIR AI*. 2023;2:e42936. doi:10.2196/42936
7. Karami M, Madloul H. Artificial intelligence and digital health in the health systems of developing countries: the challenges and vision of integration in the primary health care setting. *Front Digital Health*. 2025;7. doi:10.3389/fdgth.2025.1532361
8. Advancing medical AI with Med-Gemini. Google Research. Available from: <https://research.google/blog/advancing-medical-ai-with-med-gemini>. Accessed April 11, 2026.
9. Apple Health+ scaled back internally, new focus on incremental improvement. AppleInsider. Available from: <https://appleinsider.com/articles/26/02/06/apple-health-scaled-back-internally-will-focus-on-incremental-features-instead>. Accessed April 11, 2026.
10. Microsoft previews Copilot Health AI tool. Axios. Available from: <https://www.axios.com/2026/03/12/microsoft-copilot-health>. Accessed April 11, 2026.
11. Monlezun DJ, Omutoko L, Oduor P, et al. Digitalization of health care in low- and middle-income countries. *Bull World Health Organ*. 2025;103(2):148–154. doi:10.2471/BLT.24.291643
12. Yu J-Q, Al-Garadi MA, Guo Y, Lu Z, Smith M. Use of artificial intelligence to address health disparities in low- and middle-income countries: a thematic analysis of ethical issues. *Public Health*. 2024;234:77–83. doi:10.1016/j.puhe.2024.05.029
13. Sekalala S, Forman L, Hodgson T. Colonialism in the new digital health agenda. *BMJ Global Health*. 2024;9:e014131. doi:10.1136/bmjgh-2023-014131
14. Cazzaniga M, Jaumotte F. The global AI divide: risks of reinforcing disparities across healthcare, education, and governance. *Humanit Soc Sci Commun*. 2024. doi:10.1057/s41599-024-03947-w
15. Yang J, Soltan AAS, Clifton DA. Generalizability assessment of AI models across hospitals in a low-middle and high income country. *Nat Commun*. 2024. doi:10.1038/s41467-024-52618-6
16. Fareed M, Fatima M, Uddin J, Ahmed A, Sattar MA. A systematic review of ethical considerations of large language models in healthcare and medicine. *Front Digital Health*. 2025;7. doi:10.3389/fdgth.2025.1653631
17. Kim Y, Jeong H, Chen S, et al. Medical hallucination in foundation models and their impact on healthcare. *arXiv preprint arXiv:2503.05777* 2025. doi:10.1101/2025.02.28.25323115
18. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models. World Health Organization. Available from: <https://www.who.int/publications/i/item/9789240084759>. Accessed April 11, 2026.
19. Muralidharan V, Ng MY, AlSalamah S, et al. Global Initiative on AI for Health (GI-AI4H): strategic priorities advancing governance across the United Nations. *Npj Digital Med*. 2025;8(1). doi:10.1038/s41746-025-01618-x
20. Is giving ChatGPT health your medical records a good idea? Time. Available from: <https://time.com/7344997/chatgpt-health-medical-records-privacy-open-ai>. Accessed April 11, 2026.
21. WHO releases AI ethics and governance guidance for large multi-modal models. World Health Organization. Available from: <https://www.who.int/news/item/18-01-2024-who-releases-ai-ethics-and-governance-guidance-for-large-multi-modal-models>. Accessed April 11, 2026.
22. Cheong BC. Transparency and accountability in AI systems: safeguarding wellbeing in the age of algorithmic decision-making. *Front Human Dynamics*. 2024;6. doi:10.3389/fhumd.2024.1421273
23. Hussain SA, Bresnahan M, Zhuang J. Can artificial intelligence revolutionize healthcare in the Global South? A scoping review of opportunities and challenges. *Digital Health*. 2025;11. doi:10.1177/20552076251348024

Journal of Multidisciplinary Healthcare

Publish your work in this journal

The Journal of Multidisciplinary Healthcare is an international, peer-reviewed open-access journal that aims to represent and publish research in healthcare areas delivered by practitioners of different disciplines. This includes studies and reviews conducted by multidisciplinary teams as well as research which evaluates the results or conduct of such teams or healthcare processes in general. The journal covers a very wide range of areas and welcomes submissions from practitioners at all levels, from all over the world. The manuscript management system is completely online and includes a very quick and fair peer-review system. Visit <http://www.dovepress.com/testimonials.php> to read real quotes from published authors.

Submit your manuscript here: <https://www.dovepress.com/journal-of-multidisciplinary-healthcare-journal>

Dovepress
Taylor & Francis Group