Artificial Intelligence Adoption in HIV Public Health as a Fourth Industrial Revolution Strategy: Evidence, Economics and Governance Imperatives From Zambia

Joe Phiri[a,b*], Aaron Zimba[a], Chiyaba Njovu[a], Andrew Kashoka[c], Innocent Chiboma[c], Mulenga Chiwele[c]
[a]School of Computing, Technology and Applied Sciences, ZCAS University, Lusaka, Zambia; [b]Centre for Infectious Disease Research in Zambia (CIDRZ), Lusaka, Zambia; [c]Ministry of Health, Lusaka, Zambia

ABSTRACT

The Fourth Industrial Revolution (4IR) offers low- and middle-income countries new ways to transform public-health systems, yet artificial intelligence (AI) adoption in African health has been marked by a disconnect between technological capability and institutional readiness. This review examines the business, economic and governance dimensions of AI adoption in HIV programmes, using Zambia's national HIV electronic-health-record system as its anchoring case. Its objective is to synthesise the evidence on 4IR and health-system transformation in resource-constrained settings and to clarify the conditions under which AI adoption becomes trustworthy, equitable and sustainable. The literature was reviewed thematically across global, regional and Zambian perspectives, integrating peer-reviewed and institutional sources with empirical evidence from a multi-facility AI research programme covering 246,053 patients across six Lusaka public-health facilities. The review finds that the principal constraints on trustworthy AI are not algorithmic but structural: composite data-quality scores fall below the threshold required for reliable model training, key sociodemographic fields are largely absent from the record, and governance frameworks for responsible deployment remain underdeveloped. Interpretable models are generally preferable to complex architectures in fragmented-data settings, and data governance emerges as both the highest-value economic investment and an equity imperative, since models trained on incomplete data may act on administrative missingness rather than clinical need. The review concludes that AI adoption in public health is primarily an economic and governance challenge rather than a technological one, and that sustainable progress depends as much on investment in data governance, workforce development and regulation as on algorithmic development. It identifies specific knowledge gaps in costing, prospective evaluation and operational governance that define an agenda for future empirical research.

Article history:
Received 24 Apr 2026
Revised 10 Jun 2026
Accepted 11 June 2026
Available online 15 2026 2026

Corresponding author email:
Joe.Phiri@cidrz.org

KEYWORDS: Fourth Industrial Revolution; artificial intelligence; HIV programme management; electronic health records; digital health governance; health economics

Introduction

The Fourth Industrial Revolution (4IR) — marked by the convergence of digital, physical and biological technologies including artificial intelligence (AI), machine learning (ML), big-data analytics, the Internet of Things and cloud computing — is reshaping industries and public services worldwide (Schwab, 2016; Schwab & Davis, 2018). In healthcare, these technologies hold the prospect of precision medicine, predictive surveillance and intelligent resource allocation at a scale that was previously unachievable (Topol, 2019; Davenport & Kalakota, 2019). For sub-Saharan Africa, which carries a disproportionate burden of preventable and manageable disease while operating health systems under severe resource constraints, the 4IR is simultaneously an opportunity to leapfrog conventional development pathways and a risk of entrenching inequality where adoption lacks contextual grounding, institutional readiness and equitable governance (Nkosi, 2021; Adem et al., 2022).

Zambia's national HIV programme offers a compelling setting in which to examine 4IR health transformation.

The country operates one of the most developed HIV data ecosystems in the region, anchored by the SmartCare electronic health record (EHR) platform in use since 2004, alongside DHIS2, DATIM, DISA and ELMIS (Gumede-Moyo et al., 2019; SMART Zambia Institute, 2024). Together these systems generate millions of patient-level longitudinal records each year across hundreds of public facilities. Coupled with AI and ML analytical capacity, this infrastructure could transform programme management by anticipating which patients are at risk of tuberculosis co-infection, treatment interruption or viral-load suppression failure before adverse outcomes occur, and by directing scarce clinical resources to those most likely to benefit (Phiri, Zimba, Njovu, et al., 2026a; Phiri, Zimba, Njovu, et al., 2026b).

Yet the distance between technological potential and operational reality remains stark. Evidence from SmartCare indicates that the composite Data Quality Index across studied facilities falls below the 70% minimum threshold required for reliable AI training, and that critical sociodemographic fields — education level and household income — are effectively absent from the digital record, creating algorithmic fairness risks where tools may flag patients on the basis of administrative data gaps rather than genuine clinical need (Phiri, Zimba, Njovu, et al., 2026a; Obermeyer et al., 2019). These are not primarily technological problems; they are problems of incentive design, workforce training, data-entry culture, system configuration and regulatory oversight (Vayena et al., 2018; World Health Organization, 2021).

This article reviews the business, economic and governance dimensions of AI adoption in public-health programmes, using Zambia's HIV programme as its anchoring case. These dimensions are rarely the focus of the technical AI literature, yet they are arguably the strongest determinants of whether AI delivers sustainable, equitable value. The objectives of the review are threefold: to synthesise the evidence on 4IR and health-system transformation in low- and middle-income countries (LMICs); to characterise the economic and governance conditions under which AI adoption in HIV programmes becomes trustworthy and sustainable; and to identify the knowledge gaps that constrain responsible deployment. The remainder of the article is organised as follows. Section 2 describes the review methodology. Section 3 presents the thematic literature review across global, regional and Zambian perspectives. Section 4 sets out the theoretical framework. Section 5 presents the conceptual framework. Section 6 synthesises the literature and states the identified knowledge gaps, and Section 7 concludes.

Review Methodology

This article is a thematic, narrative review of the literature on 4IR technologies, AI economics and digital-health governance as they apply to HIV programme management in LMICs, integrated with empirical evidence from a multi-facility AI research programme conducted within Zambia's national HIV EHR system. Sources were identified through structured searches of Google Scholar, PubMed and Scopus, supplemented by institutional and grey-literature sources including World Health Organization guidance, UNAIDS fact sheets and Zambian Ministry of Health reports. Search terms combined the concepts of "artificial intelligence", "machine learning", "electronic health records", "digital health governance", "Fourth Industrial Revolution" and "HIV" with "sub-Saharan Africa", "low- and middle-income countries" and "Zambia".

The review prioritised peer-reviewed sources published between 2016 and 2026, while retaining seminal earlier work on health-economics and algorithmic-bias where it remained foundational. Inclusion criteria were relevance to AI or 4IR adoption in health systems, attention to economic or governance dimensions, and applicability to resource-constrained settings; purely technical model-development papers without deployment, economic or governance content were excluded except where they established the performance baseline relevant to LMIC contexts. The empirical evidence underpinning the Zambian case is drawn from three inter-related companion studies covering 246,053 patients across six public health facilities in Lusaka, which provide documented findings on data quality, model performance and deployment challenges. The review is therefore positioned as a synthesis that situates this primary evidence within the wider scholarly and policy literature rather than as a systematic review with formal meta-analysis.

Literature Review

3.1 The Fourth Industrial Revolution and Health-System Transformation

Schwab's (2016) conceptualisation of the 4IR identifies three technology clusters driving transformation: physical systems such as advanced robotics and 3D printing; digital systems such as AI, big data and the Internet of Things; and biological systems such as genomics and synthetic biology. In health systems the most immediately consequential cluster is digital — particularly AI and ML applied to health data — which enables automated pattern recognition, predictive modelling and decision support at a scale impossible through human analysis alone (Topol, 2019; Davenport & Kalakota, 2019).

A consistent finding across the literature is that, for LMICs, the most tractable and impactful 4IR health technology is not necessarily the most advanced. In resource-constrained settings with fragmented data infrastructure, interpretable ML models such as logistic regression and gradient boosting frequently match or outperform complex deep-learning architectures in practical deployment because they require less data and compute and produce more clinically transparent outputs (Christodoulou et al., 2019; Kelly et al., 2019). The economic implication is important: LMIC health systems do not need cutting-edge AI infrastructure to realise the value of 4IR analytics; they need the data-quality foundations and governance structures that make any ML approach reliable and equitable.

3.2 The Economic and Business Case for AI in HIV Programme Management

The economic logic of AI in HIV programme management rests on three mechanisms identified in the health-economics and digital-health literature. The first is efficiency through targeted resource allocation: risk stratification can direct clinical attention to patients most likely to experience adverse outcomes, reducing the waste of scarce resources on low-risk patients. The second is effectiveness through earlier intervention: identifying patients at risk of treatment interruption or viral-load rebound months in advance enables a shift from costly crisis response — tracing patients lost to follow-up, managing treatment failure — to preventive support. The third is equity through data-driven prioritisation: patients from marginalised backgrounds who may not self-present for care can be identified through risk profiles and proactively engaged, reducing disparities in access and quality (Vayena et al., 2018; Savedoff et al., 2012).

The literature on the cost of late-stage HIV disease management in sub-Saharan Africa supports the directional economic case. Returns from AI-guided early intervention flow primarily through three channels: averted treatment failure, which reduces the cost of second-line antiretroviral therapy and hospitalisation; reduced tuberculosis treatment costs through earlier detection; and efficiency savings from better-targeted adherence support (Mugglin et al., 2020; Gupta et al., 2013). A recurring theme is distributional: the costs of AI investment fall on governments, development partners and research institutions, while the benefits accrue to patients — disproportionately those who are socioeconomically marginalised — and to the health system through reduced cost of treatment-failure management. Crucially, the same evidence warns that models deployed without first addressing data-quality deficiencies may flag patients on the basis of administrative missingness rather than clinical need, making data governance both a technical prerequisite and an equity imperative (Obermeyer et al., 2019; Panch et al., 2019).

3.3 The African and Zambian Context: Infrastructure, Assets and Gaps

Several African countries have leveraged 4IR technologies for health-system improvement — Rwanda's drone delivery of blood products, South Africa's ML-based tuberculosis screening from chest radiographs, and Kenya's integration of mobile-health platforms into community health systems (Abutaleb, 2020; Hatherill et al., 2021). These cases show that 4IR adoption in Africa is not hypothetical, but they reveal a consistent pattern: successful adoption is preceded by substantial investment in foundational infrastructure — connectivity, digital literacy, system integration — and in governance, including regulatory frameworks, data-protection law and ethics oversight (Nkosi, 2021; Abutaleb, 2020).

Zambia's HIV programme is one of the largest in the region, with approximately 1.3 million people living with HIV as of 2024 and antiretroviral therapy available at public facilities nationwide (UNAIDS, 2024). Achievement of the UNAIDS 95–95–95 targets depends on identifying and retaining the most vulnerable patients — those at risk of treatment interruption, viral-load rebound or tuberculosis co-infection (Ministry of Health, 2022). Three clinical outcomes drive programme performance: incident tuberculosis co-infection, interruption in treatment and unsuppressed viral load, each requiring a distinct programmatic response. AI-driven risk stratification can enable managers to direct these differentiated responses to the patients most likely to benefit before adverse outcomes occur (Phiri, Zimba, Njovu, & Lumpa, 2026).

Zambia's digital-health infrastructure comprises interconnected systems — SmartCare for the national HIV EHR, DHIS2 for aggregate reporting, DISA for laboratory information, ELMIS for the pharmaceutical supply chain and DATIM for donor reporting — which together produce a comprehensive digital footprint of the programme (Gumede-Moyo et al., 2019; SMART Zambia Institute, 2024). SmartCare alone holds records for more than 246,000 people living with HIV across the six major Lusaka facilities studied, with data spanning 2003 to 2025. This longitudinal depth is a significant asset, yet the infrastructure also carries well-documented limitations: education-level completeness of only 1.2–7.9%, engagement-feature coverage for only 12.2% of patients, and a composite Data Quality Index of 53.5–58.3% — all below the 70% threshold for reliable AI training (Phiri, Zimba, Njovu, et al., 2026a). These gaps reflect data-entry practices, system-configuration choices and the absence of mandatory-field requirements — governance problems with governance solutions.

Theoretical Framework

This review is underpinned by the theory of the Fourth Industrial Revolution as articulated by Schwab (2016) and elaborated by Schwab and Davis (2018). At its core, 4IR theory holds that contemporary transformation is distinguished from earlier industrial revolutions not by any single technology but by the velocity, breadth and systemic depth with which digital, physical and biological technologies converge and reinforce one another. The theory emphasises that technological capability alone does not determine impact; outcomes are shaped by the institutional, economic and governance systems within which technologies are embedded.

This framing fits the present study for three reasons. First, it directs analytical attention away from the algorithm in isolation and toward the surrounding system of data infrastructure, workforce capacity and regulation — precisely the layers where the Zambian evidence locates the binding constraints on trustworthy AI. Second, its systemic emphasis accommodates the multi-actor reality of public-health AI, in which ministries, research institutions, development partners and facilities each hold a part of the value chain. Third, its concern with equity and the risk of entrenching inequality speaks directly to the algorithmic-fairness risks that arise when models are trained on data with systematic demographic missingness. The 4IR lens is therefore used not as a descriptive label but as an analytical commitment: that the determinants of value are institutional and economic as much as technical.

Conceptual Framework

Drawing the literature together, the conceptual framework presents AI-enabled HIV programme transformation as a layered system in which analytical capability is valuable only to the extent that it rests on sound data governance beneath it and feeds clinical decision support and institutional oversight above it. Figure 1 sets out the framework as six interdependent layers, from the digital-infrastructure foundation through data governance and the AI/ML analytical engine to clinical decision support, governance and ethics, and ultimately health outcomes and impact. Each layer is necessary but insufficient without the others, which is the central proposition the framework expresses.

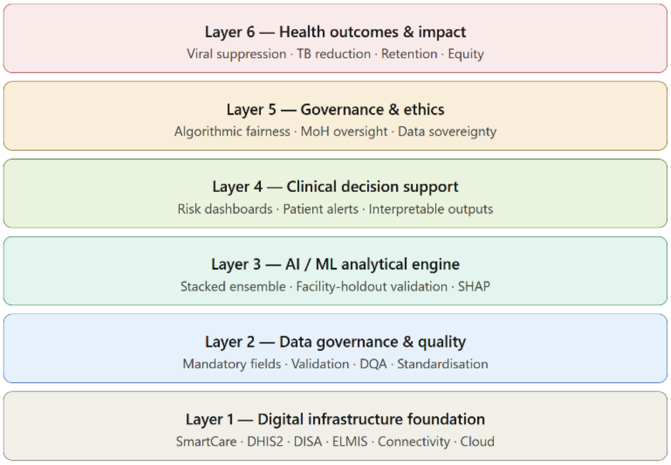

Figure 1. 4IR Technology Framework for HIV Programme Transformation in Zambia. Six interdependent layers from digital infrastructure to health outcomes; the AI/ML analytical engine (Layer 3) requires data-governance foundations (Layers 1–2) and clinical decision support and governance (Layers 4–6) to deliver sustainable impact.

The framework also expresses a value chain. Figure 2 traces value from raw EHR data through ML processing, clinical output and health-worker action to patient outcomes and the achievement of the Sustainable Development Goals. Each link is simultaneously a technical requirement and an economic investment decision: data-input quality determines prediction reliability; output interpretability determines health-worker uptake; and the effectiveness of health-worker action determines the ultimate impact on patient outcomes.

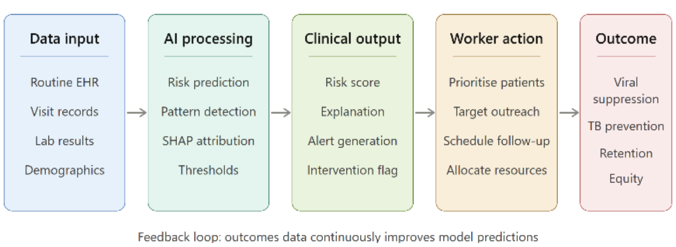

Figure 2. AI Value-Creation Chain in HIV Programme Management. Value flows from data input through ML processing, clinical decision support and health-worker action to patient outcomes; each link represents both a technical and an economic investment decision.

The economics of the framework are illustrated in Figure 3, which shows an indicative distribution of implementation cost and a five-year cost–benefit trajectory. The framework assigns the largest single share of implementation cost to data governance rather than to algorithmic development, reflecting the literature's consistent finding that data quality, not model sophistication, is the binding constraint on reliable AI in this context. This yields a counter-intuitive but actionable proposition: investment in data governance — such as making sociodemographic fields mandatory in SmartCare — is the highest-value individual intervention for improving AI reliability and economic return.

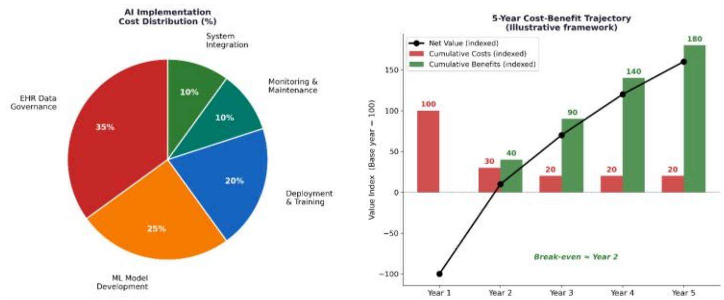

Figure 3. Economic Analysis of AI Investment in HIV Programme Management. Left: indicative distribution of implementation cost across categories. Right: indicative five-year cost–benefit trajectory showing break-even around Year 2 and sustained net positive value from Year 3 onward.

Synthesis and Knowledge Gaps

Synthesising across the themes reviewed, three points of broad agreement emerge. First, the literature converges on the view that the principal obstacles to trustworthy AI in LMIC health systems are not algorithmic but structural — data-quality deficiencies, institutional-capacity gaps and absent governance frameworks. Second, there is consensus that interpretable models are generally preferable to complex architectures in resource-constrained, data-fragmented settings. Third, the economic and equity literatures agree that the distribution of costs and benefits, and the risk of acting on administrative missingness rather than clinical need, make data governance an equity concern and not merely a technical one. A genuine tension also runs through the literature: the same AI tools that can advance equity by surfacing under-served patients can entrench inequality if deployed on poor-quality data, and the field has not resolved how to reconcile rapid adoption with the slower work of building data and governance foundations.

Effective governance responses to these tensions can be organised around four design principles drawn from the literature — accountability, transparency, equity and sustainability (Vayena et al., 2018; Panch et al., 2019;

World Health Organization, 2021). In the Zambian programme these principles are operationalised through an institutional architecture in which the Ministry of Health occupies the central governance-authority role, supported by complementary functions from academia, research institutions, multilateral and bilateral partners, the ICT sector, health facilities and an ethics-review body. Figure 4 presents this architecture, and Figure 5 sets out a phased adoption roadmap through which the governance and technical capabilities can be built sequentially rather than assumed.

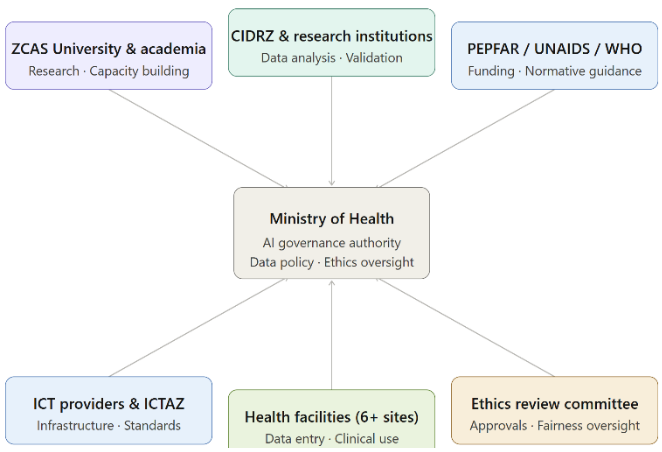

Figure 4. Institutional Governance Architecture for AI in Zambia's HIV Programme. The Ministry of Health occupies the central governance-authority role, with six institutional actors contributing complementary governance functions.

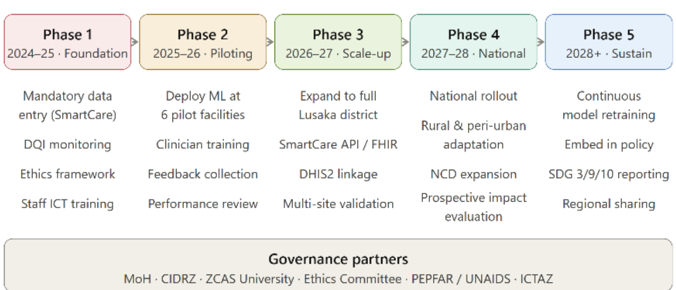

Figure 5. Five-Phase AI Adoption Roadmap for Zambia's HIV Programme (2024–2030). Each phase builds governance and technical capability sequentially, with concrete milestones and institutional responsibilities.

Against this synthesis, the review identifies four specific knowledge gaps. First, the economic evidence remains directional rather than quantified: no facility-level costing study has yet measured the cost and return of AI-guided intervention in an LMIC HIV programme, leaving the business case argued rather than demonstrated. Second, there is an absence of prospective deployment evidence — the literature documents model performance and data-quality conditions but not whether AI-guided intervention changes patient outcomes in routine clinical workflows. Third, the operational governance literature is largely conceptual; widely cited guidance such as the World Health Organization's (2021) framework specifies principles but offers limited evidence on which concrete governance mechanisms — mandatory-field configuration, missingness-confounding testing, fairness assessment — are effective and feasible in practice. Fourth, the equity literature has established that demographic missingness can bias algorithms, but has not established how missingness-aware deployment protocols perform when implemented at scale in a national programme. These gaps are specific and logically derived from the evidence reviewed, and they define a concrete agenda for empirical research.

Conclusion

This review has argued that trustworthy AI adoption in public health is primarily an economic and governance challenge rather than a technological one. The literature, read alongside evidence from Zambia's HIV programme, consistently locates the binding constraints not in algorithmic sophistication but in data quality, institutional capacity and governance. The contribution of this review is to integrate the technical, economic, equity and governance literatures into a single layered framework, and to show that the highest-value investment in AI for health is frequently the least technologically glamorous — the data-governance foundation on which any reliable, equitable model must stand.

The review identifies a clear agenda for future empirical research: facility-level costing studies to quantify the economic case; prospective evaluations to measure whether AI-guided intervention changes patient outcomes; implementation studies to test which concrete governance mechanisms are effective and feasible; and equity-focused evaluations of missingness-aware deployment at programme scale. For Zambia and comparable LMICs, the HIV programme represents the most mature and impactful starting point for digital-health transformation. Realising that opportunity will depend less on acquiring advanced algorithms than on the disciplined work of building data governance, transparent and equitable deployment, and sustained institutional commitment — foundations that can, in time, extend AI-enabled transformation well beyond HIV to the wider spectrum of the country's health challenges.

Conflict of Interest

The authors declare that this research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. The institutional affiliation of two co-authors with ZCAS University is acknowledged and disclosed.

References

Abutaleb, B. (2020). Rwanda's drone delivery system: A model for Africa? The Lancet Digital Health, 2(4), e156–e157. https://doi.org/10.1016/S2589-7500(20)30050-3

Adem, A., Yusuf, A., & Mwangi, J. (2022). Digital health infrastructure in sub-Saharan Africa: Opportunities and constraints. Health Policy and Planning, 37(3), 372–384. https://doi.org/10.1093/heapol/czab140

Christodoulou, E., Ma, J., Collins, G. S., Steyerberg, E. W., Verbakel, J. Y., & Van Calster, B. (2019). A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. Journal of Clinical Epidemiology, 110, 12–22. https://doi.org/10.1016/j.jclinepi.2019.02.004

Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. Future Healthcare Journal, 6(2), 94–98. https://doi.org/10.7861/futurehosp.6-2-94

Gumede-Moyo, S., Todd, J., Bond, V., Mee, P., & Filteau, S. (2019). A qualitative inquiry into implementing SmartCare for PMTCT data in Zambia. BMJ Open, 9(9), e030428. https://doi.org/10.1136/bmjopen-2019-030428

Gupta, R. K., Lawn, S. D., Bekker, L.-G., Wood, R., & Pillay, D. (2013). Impact of HIV and CD4 count on tuberculosis diagnosis. The International Journal of Tuberculosis and Lung Disease, 17(8), 1014–1022. https://doi.org/10.5588/ijtld.13.0143

Hatherill, L., Scriba, T. J., & Penn-Nicholson, A. (2021). AI diagnostics for tuberculosis in Africa: Progress and challenges. Bulletin of the World Health Organization, 99(10), 738–745. https://doi.org/10.2471/BLT.20.279307

Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. BMC Medicine, 17, 195. https://doi.org/10.1186/s12916-019-1426-2

Ministry of Health, Zambia. (2022). Zambia consolidated guidelines for treatment and prevention of HIV infection. Ministry of Health.

Mugglin, C., Kläger, D., Gueler, A., Vanobberghen, F., Rice, B., & Egger, M. (2020). Long-term retention on antiretroviral therapy among infants, children, adolescents and adults in low- and middle-income countries. Clinical Infectious Diseases, 71(5), e113–e121. https://doi.org/10.1093/cid/ciz1213

Nkosi, M. (2021). 4IR in Africa: Opportunities, challenges and policy responses. Development Southern Africa, 38(1), 1–15. https://doi.org/10.1080/0376835X.2020.1837599

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447–453. https://doi.org/10.1126/science.aax2342

Panch, T., Mattie, H., & Atun, R. (2019). Artificial intelligence and algorithmic bias: Implications for health systems. Journal of Global Health, 9(2), 010318. https://doi.org/10.7189/jogh.09.020318

Phiri, J., Zimba, A., Njovu, C., Kashoka, A., Chiboma, I., & Chiwele, M. (2026a). Trustworthy AI in HIV public health: A multi-facility data quality assessment of electronic health records in Zambia [Manuscript submitted for publication]. Discover Artificial Intelligence.

Phiri, J., Zimba, A., Njovu, C., Kashoka, A., Chiboma, I., & Chiwele, M. (2026b). An explainable multi-outcome machine learning framework for personalised clinical decision support in HIV care using routine electronic health records in Zambia [Manuscript submitted for publication]. BMC Medical Informatics and Decision Making.

Phiri, J., Zimba, A., Njovu, C., & Lumpa, M. (2026). A systematic review of AI and ML methods and deployment challenges for public-health predictions using electronic health records in low- and middle-income countries [Manuscript under revision]. Discover Artificial Intelligence.

Savedoff, W. D., de Ferranti, D., Smith, A. L., & Fan, V. (2012). Political and economic aspects of the transition to universal health coverage. The Lancet, 380(9845), 924–932. https://doi.org/10.1016/S0140-6736(12)61083-6

Schwab, K. (2016). The Fourth Industrial Revolution. World Economic Forum.

Schwab, K., & Davis, N. (2018). Shaping the future of the Fourth Industrial Revolution. World Economic Forum.

SMART Zambia Institute. (2024). SmartCare Pro — Advancing healthcare. https://www.szi.gov.zm

Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. Nature Medicine, 25, 44–56. https://doi.org/10.1038/s41591-018-0300-7

UNAIDS. (2024). Global HIV & AIDS statistics — Fact sheet. https://www.unaids.org/en/resources/fact-sheet

Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. PLoS Medicine, 15(11), e1002689. https://doi.org/10.1371/journal.pmed.1002689

World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. World Health Organization.