Artificial Intelligence-Enabled Smart Systems for Sustainable Healthcare and Environmental Monitoring: Socio-Economic Perspectives

Alvena Shahid[68], Roberto Acevedo[69]

https://doi.org/10.66793/intraders15proceeding17

Abstract

Artificial intelligence has moved well beyond computer science. It now shapes healthcare delivery, environmental governance, and economic planning in ways that were difficult to predict even a decade ago. This chapter examines how AI-driven smart systems are changing these domains, with particular focus on developing countries, where the gap between what the technology can do and what actually gets deployed remains frustratingly wide. Drawing on literature across engineering, public health, and development economics, the chapter covers three things: how intelligent monitoring systems are built, what they do in clinical and environmental settings, and what their broader socio-economic consequences look like. The evidence is mixed. AI does deliver real gains in diagnostic accuracy, environmental surveillance, and resource efficiency. But those gains are unevenly distributed. Data infrastructure deficits, algorithmic bias, and chronic shortages of technical expertise all hit hardest in the places that need this technology most. The chapter also proposes a framework for evaluating AI readiness in resource-limited settings, and identifies where targeted investment and policy reform could actually move the needle. Pakistan is used throughout as a reference case, not because it is uniquely representative, but because it makes the tradeoffs concrete.

Keywords: artificial intelligence, smart systems, healthcare monitoring, environmental surveillance, sustainable development, developing countries, socio-economic impact

1. Introduction

New technologies do not automatically produce better outcomes [1]. This is probably the most important thing to say upfront, and it is also the thing most often glossed over in discussions of AI [2].

What tools actually do is amplify existing capacities, expose existing weaknesses, and generate new categories of risk alongside their benefits [3]. AI is no different. What makes it worth taking seriously is scale and speed. Algorithms now screen medical images with accuracy that approaches, and in specific tasks exceeds, that of trained specialists. Networks of low-cost sensors generate environmental data streams that no human workforce could process manually [4]. Predictive models flag equipment failures days before they happen. None of this is hypothetical. These systems are running right now.

[68] Facultad de Ingeniería. Universidad San Sebastián. Bellavista 7. 8420524. Santiago. Chile Email: alvena.shahid@gmail.com, ORCID: https://orcid.org/0000-0001-6557-8207
[69] Facultad de Ingeniería. Universidad San Sebastián. Bellavista 7. 8420524. Santiago. Chile Email: roberto.acevedo.llanos@gmail.com, ORCID: 0000-0001-6847-0285 (Corresponding author)

But who benefits? That is the harder question. A diagnostic algorithm trained on data from high-income populations may perform poorly on patients from South Asia or Sub-Saharan Africa [5]. A smart city platform built for reliable broadband cannot function where internet access cuts out regularly [6-11]. The socio-economic implications of AI are therefore inseparable from questions of data governance, institutional capacity, and political economy [12-18].

This chapter takes those questions seriously. It begins with a technical overview of AI systems, then works through healthcare and environmental applications, and turns finally to economic and social dimensions. Section 5 presents a framework for assessing AI readiness in developing-country contexts. The chapter closes with governance challenges and some honest observations about what would actually need to change for these benefits to be more widely shared.

Pakistan recurs throughout as a reference case. It is not uniquely representative. But it illustrates, with considerable clarity, both the real promise and the real constraints of AI adoption in resource-limited settings.

2. Artificial intelligence in engineering and smart systems

2.1 Foundations of artificial intelligence

In its current form, AI refers to computational systems that process data, identify patterns, and generate outputs, whether predictions, classifications, recommendations, or control actions, without being explicitly programmed for each task. The term covers a wide spectrum. The approaches differ significantly in architecture, data requirements, and interpretability.

Machine learning (ML) is the broadest category. ML systems learn from training data by adjusting internal parameters to minimize prediction error. Classical ML algorithms, including support vector machines, random forests, and gradient boosting, work well for structured tabular data and run on modest hardware. Deep learning (DL) is a subset of ML. It uses layered neural networks to extract hierarchical representations from raw inputs like images, audio, or text. DL drove most of the headline advances in medical imaging and speech recognition over the past decade. It also requires far more data and computing power than classical ML, and its internal workings are considerably harder to interpret.

Reinforcement learning, a third paradigm, trains agents through trial-and-error interaction with an environment. It optimizes for long-term cumulative reward. Autonomous systems, robotic control, and dynamic resource allocation all draw on this approach. In practice, most real AI systems combine elements of several paradigms rather than fitting neatly into one.

Figure 1 shows the hierarchical taxonomy of these AI approaches and maps them onto representative application domains discussed throughout this chapter.

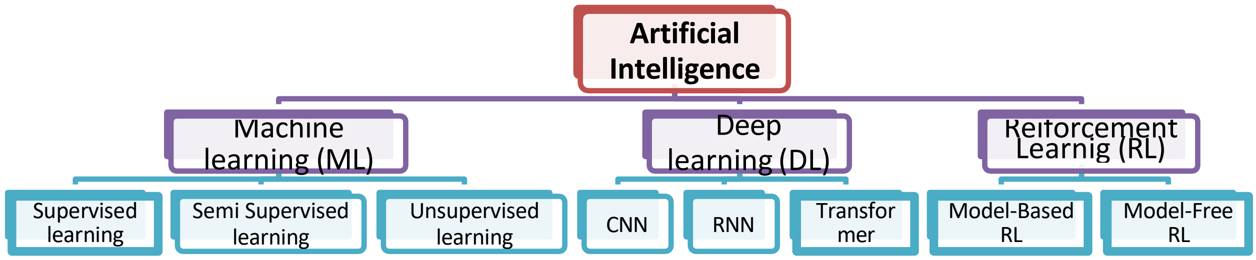

Figure 1: Taxonomy of Artificial Intelligence Approaches and Application Domains

Table 1 below shows a structured comparison of the three main AI paradigms, including their data needs, example algorithms, and areas of application that are important to this chapter.

Table 1: Comparison of Principal Artificial Intelligence Paradigms

AI Paradigm	Data Requirements	Representative Algorithms	Compute Demand	Example Applications
Machine Learning (ML)	Labeled / unlabeled structured data	Support Vector Machine, Random Forest, Gradient Boosting	Low–medium	Clinical risk stratification, crop yield prediction
Deep Learning (DL)	Large labeled datasets; images, audio, text	CNN, RNN, LSTM, Transformer	High	Medical imaging, NLP, speech recognition
Reinforcement Learning (RL)	Interaction with environment; reward signals	Q-Learning, Actor-Critic, PPO	High	Autonomous systems, smart grid control
Hybrid /	Combination of	XGBoost + NN,	Medium–high	Wearable

Ensemble	above	federated models		health monitoring, satellite analysis

Source: Compiled from Russell & Norvig (2020) and domain literature.

2.2 Characteristics of smart systems

A "smart system" has three defining properties. It acquires data continuously from its environment. It processes that data to infer current or future states. And it acts on those inferences, either autonomously or by giving human operators the information they need to decide.

The key difference from conventional automation is adaptability. A traditional control system follows fixed rules. A smart system updates its behavior as conditions change. That adaptability matters in clinical settings, where patient conditions shift unpredictably, and in ecological monitoring, where pollutant levels respond to weather, traffic, and industrial activity in real time.

The infrastructure underlying these systems typically includes sensor networks for data acquisition, communication protocols for transmission, cloud or edge computing for processing, and user interfaces for presenting results. Each layer has its own requirements and its own failure modes.

2.3 Convergence of AI with the Internet of Things

The convergence of AI with the Internet of Things (IoT) is probably the most consequential development in this area. IoT devices, physical objects embedded with sensors, actuators, and connectivity, generate continuous streams of structured, time-stamped data that are well-suited to ML analysis. AI, in turn, allows IoT networks to move beyond passive data collection toward active inference and control.

Figure 2: Architecture of an AI-Enabled Smart Monitoring System

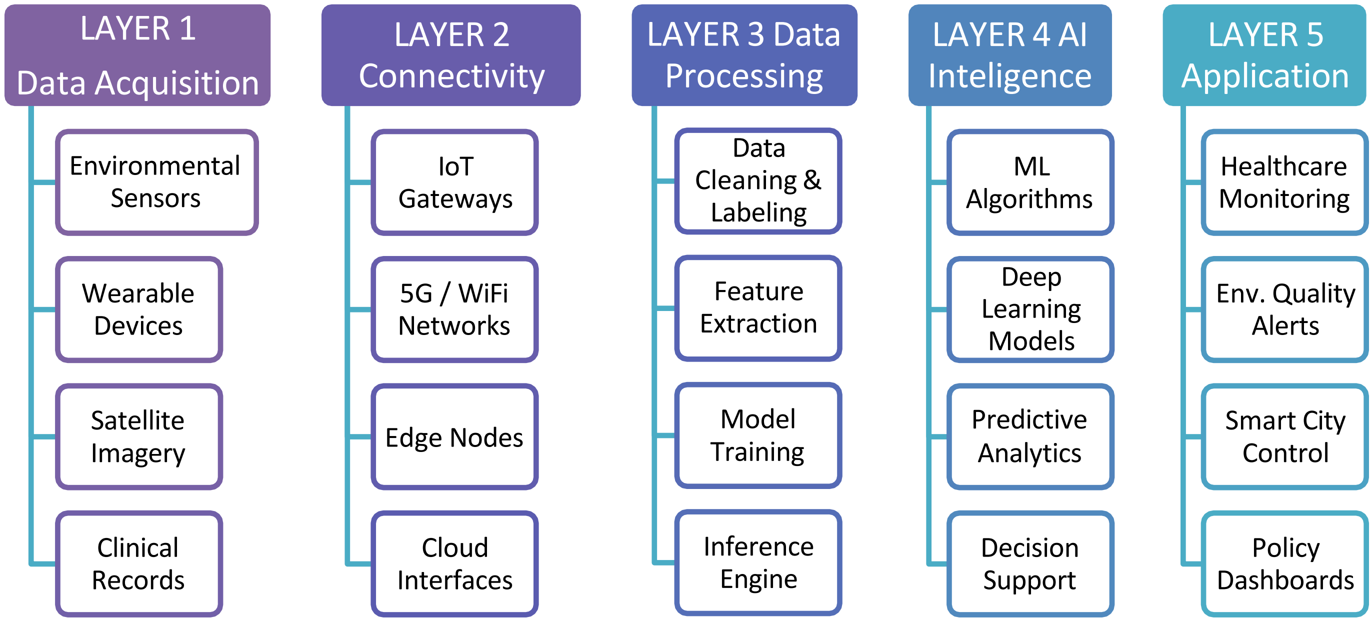

Figure 2: Architecture of an AI-Enabled Smart Monitoring System (Five-Layer Model)

In healthcare, this shows up in wearable monitors that detect arrhythmias, alert caregivers to deteriorating vital signs, or track medication adherence. In environmental monitoring, it appears in low-cost air quality sensor networks that achieve spatial resolution far beyond what traditional fixed stations can manage. Researchers call the result "cyber-physical systems," tightly coupled integrations of digital intelligence and physical sensing that can operate at scales and speeds no human-operated system could match (Gubbi et al., 2013; Shi et al., 2016).

3. AI in sustainable healthcare systems

3.1 Intelligent health monitoring and wearable technologies

Wearable health technology has grown substantially over the past decade. Better sensors, wireless communication, and battery efficiency all contributed. Today's wearables can continuously monitor heart rate, oxygen saturation, ECG morphology, skin temperature, galvanic skin response, and movement patterns. When combined with AI-based signal processing, these streams yield clinically meaningful inferences that go well beyond what the raw numbers suggest.

Photoplethysmography-based wearables can now detect atrial fibrillation with sensitivity and specificity comparable to standard ECG screening in undiagnosed populations (Turakhia et al., 2019). Continuous glucose monitors paired with machine learning can anticipate hypoglycemic episodes 30 to 60 minutes before they occur, giving Type 1 diabetes patients time to act (Oviedo et al., 2017). These are genuine clinical capabilities. The honest caveat is that the devices remain expensive relative to household incomes in low- and middle-income countries, and the algorithms have typically been validated on populations that differ demographically from users in developing-country contexts. That gap matters.

3.2 Remote healthcare and AI-augmented telemedicine

Telemedicine has existed in various forms for decades. What changed during the COVID-19 pandemic was the pace of adoption, driven by necessity when in-person care was disrupted globally. AI adds real value to telemedicine platforms. Natural language processing models can transcribe and structure clinical encounters in real time. AI triage systems assess symptom severity before a consultation begins. Computer vision tools can analyze patient-submitted photographs, skin lesion images, wound progress photos, fundus images, with preliminary accuracy that holds up for screening purposes.

In Pakistan, where the physician-to-population ratio is approximately 1:1,000 against a WHO recommendation of 1:600, and where rural areas face far more acute shortages than that national average implies, this matters practically (Pakistan Medical Commission, 2022). A dermatologist reviewing AI-pre-screened images can work through more cases per hour than one doing unassisted review. That is not a transformative claim. It is a specific, achievable one.

3.3 Predictive diagnostics and disease surveillance

Predictive diagnostics uses historical patient data to estimate the probability of future disease events. It is one of the areas where AI has produced the most credible clinical results. Algorithms trained on electronic health records can flag patients at elevated risk of sepsis, acute kidney injury, or hospital readmission hours or days before deterioration becomes apparent, enabling early intervention (Rajpurkar et al., 2022).

Medical imaging is where the evidence base is deepest. Convolutional neural networks now match or exceed radiologist performance on specific tasks: detecting diabetic retinopathy from fundus photographs, identifying malignant pulmonary nodules on chest CT, finding breast cancer on mammography. These results are task-specific. They do not transfer automatically to adjacent problems, and performance is sensitive to how closely training data matches the deployment population.

For disease surveillance, AI enables population-level early warning by integrating data from clinical records, pharmacy claims, lab results, mortality registries, and search query patterns. Monitoring platforms can detect anomalous signals indicative of emerging outbreaks days before traditional surveillance channels flag them. In countries with limited formal surveillance infrastructure, that lead time is genuinely valuable.

Table 2 gives a summary of the evidence base for some AI healthcare applications. It shows both the results that have been shown and the problems that make it hard to apply them to LMIC settings.

Table 2: AI Applications in Healthcare - Evidence Summary and LMIC Limitations

Application	Technology Used	Evidence / Outcome	Limitation / Caveat
Arrhythmia Detection	Wearable PPG + DL classifier	Sensitivity / specificity comparable to ECG screening (Turakhia et al., 2019)	Device cost; skin-tone bias in PPG sensors
Diabetic Retinopathy Screening	Fundus photography + CNN	AUC > 0.97 in multiple validation studies; FDA-cleared system (IDx-DR)	Training data predominantly from high-income populations
Sepsis Prediction	EHR-integrated ML (LSTM)	12–24 h early alert; 20% reduction in ICU mortality in prospective trials	Requires real-time EHR integration; alert fatigue risk
Hypoglycaemia Prediction	Continuous glucose monitor + ML	30–60 min advance warning; reduces severe episodes by ~40% (Oviedo et al., 2017)	Device affordability in LMICs; connectivity dependency
COVID-19 Outbreak Detection	Multistream AI surveillance (HealthMap)	3–7 day earlier detection vs. official reporting (Brownstein et al., 2020)	Mis-triage risk; social media noise
Radiology Triage	Chest X-ray CNN (CheXNet,	Expert-level TB and pneumonia	Local dataset scarcity; regulatory approval

	CheXpert)	detection; reduces reporting backlog	timelines

Source: *Synthesized from cited literature. AUC = Area Under the Receiver Operating Curve; PPG = Photoplethysmography; EHR = Electronic Health Record.*

4. Environmental monitoring and AI-enabled smart solutions

4.1 Air and water quality monitoring

Environmental monitoring has historically been constrained by two things: the high cost of regulatory-grade instruments and the labor required for manual sampling. Both constraints are easing as low-cost sensor networks and AI-based data processing become more accessible.

Air quality monitoring shows the tradeoffs clearly. Regulatory monitoring stations are accurate but sparse. A typical city has four to eight of them, producing a spatial average that may say very little about actual exposure near industrial zones or high-traffic corridors. Low-cost electrochemical and optical sensors allow networks of hundreds of nodes. But they drift, and without correction they can produce substantial measurement error (Snyder et al., 2013). ML models trained on co-located measurements from low-cost and reference instruments can correct sensor-specific biases and extend those corrections across the network. Lahore, which ranks among the most polluted cities globally by PM2.5 concentration, has begun exploring hybrid approaches rather than relying exclusively on its small number of reference stations.

Water quality monitoring follows similar logic. AI-enhanced sensor systems can detect contaminants, heavy metals, nitrates, biological indicators, in real time rather than waiting days for laboratory results. In agricultural settings, AI models integrating satellite imagery, soil sensor data, and weather inputs can predict irrigation needs and fertilizer application rates that reduce nutrient runoff.

4.2 Climate data analysis and predictive modeling

Climate science is one domain where AI is changing what is computationally feasible. Traditional numerical weather prediction models solve physical equations governing atmospheric dynamics on spatial grids. They work at global to regional scales but are expensive and poorly resolved at the local level.

Deep learning approaches are proving useful for complementary tasks: downscaling coarse climate model outputs to local resolution, correcting model projections against observed data, and identifying precursors of extreme events in large historical datasets. Google DeepMind's GraphCast architecture has demonstrated AI-based weather forecasting with accuracy comparable to the best operational numerical models at a fraction of the computational cost. That has real implications for countries without access to large high-performance computing infrastructure (Lam et al., 2023).

For Pakistan, which faces accelerating glacier melt, shifting monsoon patterns, and more frequent severe floods, better local climate projections translate into better planning for water resources, agriculture, and disaster risk management. The link between the science and the policy is direct.

4.3 Smart cities and urban system integration

Smart city initiatives use networked sensors, data analytics, and AI-based decision support to improve urban services. AI-based traffic signal optimization has consistently reduced average travel time and vehicle emissions by adapting signal timing to actual traffic conditions. Predictive maintenance systems applied to water distribution infrastructure use sensor data and ML models to identify pipes approaching failure before they rupture, cutting both service disruption and emergency repair costs.

For cities in developing countries, the "smart city" concept needs careful adaptation. Wholesale adoption of platforms designed for high-income infrastructure conditions is rarely feasible. A more productive approach is modular. Identify specific bottlenecks where sensor data and AI decision support can produce measurable improvement. Implement at manageable scale. Build local technical capacity alongside the technology rather than after it.

5. Conceptual framework: evaluating AI readiness in developing countries

Assessing whether an AI application is appropriate for a developing-country context requires more than asking if the technology works. The real question is whether the full system, human, institutional, and technical components together, can function reliably under local conditions. Table 3 shows the following framework, which lists four areas that are worth looking into.

Table 3: AI Readiness Framework Dimensions, Indicators, and Priority Interventions (Pakistan Context)

Readiness Dimension	Key Indicators	Pakistan Context	Priority Interventions
Technical Infrastructure	Internet penetration, broadband speed, power reliability, cloud access	Pakistan: ~47% internet penetration; frequent load-shedding constrains	Edge AI deployment; offline-first architectures

		cloud-dependent systems	
Data Ecosystem	Availability of labeled local datasets, data standards, interoperability, data governance	Limited digitized health records; fragmented environmental sensor networks	National data-sharing frameworks; open government data policies
Institutional & Human Capacity	AI researchers per capita, STEM graduates, hospital IT readiness, government digitization	~25,000 IT graduates/year (PSEB, 2022); shortage of ML specialists in public sector	AI curriculum in universities; public-sector fellowships
Regulatory & Governance Frameworks	Data protection law, AI liability frameworks, algorithmic audit requirements	Personal Data Protection Act (2023) in place; AI-specific regulation still nascent	Dedicated AI regulatory agency; mandatory bias audits for health AI
Workforce Digital Literacy	End-user AI literacy, clinician digital competency, public awareness	High variation across urban/rural divide; low AI literacy among clinical staff	Continuing professional education; community health worker training
Affordability & Access	Device cost relative to income, subsidized access programs, open-source availability	Wearables and diagnostic devices unaffordable for most; few public subsidy schemes	Government procurement programs; open-source AI model libraries

Source: *Authors' framework, informed by ITU Digital Readiness Index and World Bank (2023).*

Technical infrastructure readiness covers connectivity (bandwidth, reliability, geographic coverage), computing infrastructure (cloud access, local processing capacity, power reliability), and sensor network maturity. Applications that require real-time data processing over unreliable connections need offline functionality or edge-computing capability built in from the start, not added as an afterthought.

Data ecosystem readiness covers locally relevant training data, standards for data collection and labeling, interoperability between data systems, and governance frameworks for data quality and security. AI models trained on data from different populations may not generalize. The absence of local labeled datasets is a genuine constraint, not a minor technical detail.

Institutional and human capacity include ... critically evaluate AI systems, and the organizational structu ... agnostic tool deployed in a clinic without staff trained to i n ... will not improve patient outcomes. The tool is not the interv

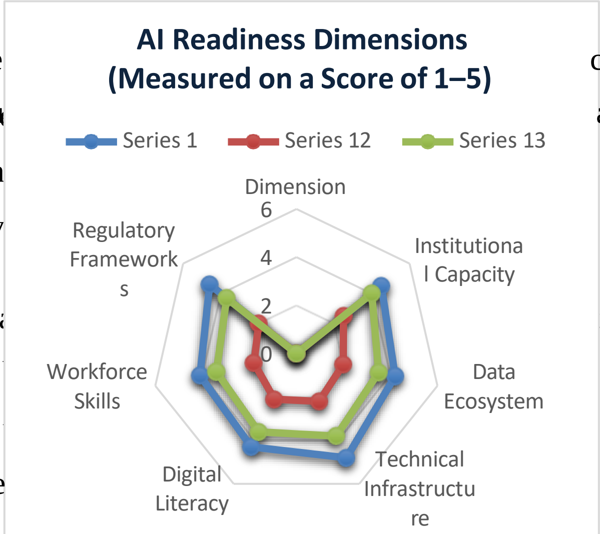

Regulatory and governance frameworks ... ructures within which AI systems operate. Who is responsi ... How is patient data protected? How is algorithmic bias de ... are still building the regulatory capacity to answer these que

Figure 3 shows the readiness profile as a radar chart. It compares the average scores of developed economies to the estimated current and target scores of Pakistan and other LMICs. The difference between the current and target profiles sets the agenda for the intervention. Before AI can reliably deliver on its promises, each of these areas needs targeted investment: connectivity, data infrastructure, regulatory development, and workforce capacity.

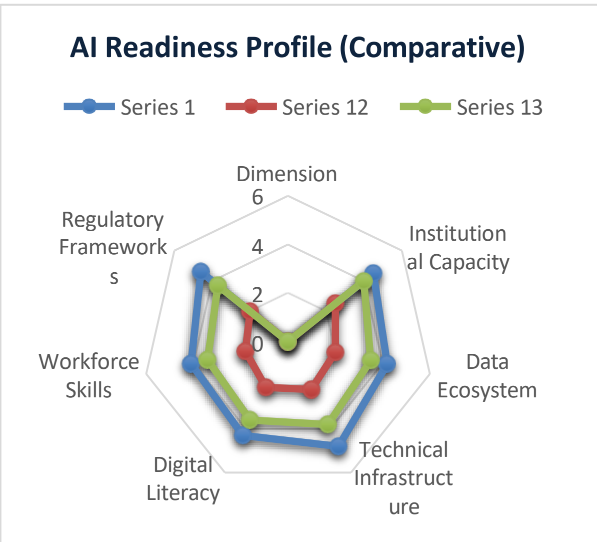

Figure 3: AI Readiness Framework: Comparative Profile for Developed vs. Developing Economies (Score 1–5)

This framework is not a checklist. It is an analytical tool for identifying where intervention is needed, whether that means improving connectivity, collecting better data, running training programs, or drafting regulatory guidance, before or alongside technology deployment.

6. Economic and social impacts of AI

6.1 AI and economic development

The macroeconomic literature on AI and growth is still forming. The technology is recent enough that long-run effects are not yet visible in the data. What is clearer from firm-level and sector-level studies is that AI adoption tends to improve productivity within the adopting unit. The magnitude varies considerably by sector, task type, and implementation quality.

For developing economies, the relevant question is whether AI can support catch-up by enabling efficiency gains in sectors that currently operate well below international best practice. Agriculture, healthcare, and logistics look credible. AI-based crop advisory systems, AI-augmented primary care, and AI-optimized supply chains all address bottlenecks with tangible economic costs.

A more cautious view, developed by economists including Dani Rodrik, is that AI may actually compress opportunities for late industrializers. If AI enables automation of tasks that previously required cheap labor, the comparative advantage that drove earlier manufacturing-led development may be harder to exploit today (Rodrik, 2018). This is not a settled debate, but it is a serious one that development planners should engage with rather than ignore.

6.2 Social transformation: healthcare access, education, and quality of life

In healthcare, the argument for AI in developing countries is not that it replaces specialist care. It is that it extends the reach of the health workforce that exists. A community health worker equipped with a validated AI screening tool for childhood malnutrition, tuberculosis, or diabetic retinopathy can identify cases that would otherwise go undetected until they become much more expensive to treat.

In education, AI-powered adaptive learning systems adjust instructional content to individual student performance. This potentially addresses a persistent problem: heterogeneous classrooms where teachers cannot simultaneously serve students at very different ability levels. Evidence from controlled trials in several developing-country contexts shows modest but real learning gains from well-implemented adaptive tools, particularly in mathematics (Muralidharan et al., 2019). Modest but real is worth taking seriously.

6.3 The distributional problem

The benefits above are real. They are also not guaranteed, and they do not distribute themselves. Technologies adopted primarily by better-resourced urban hospitals, wealthier farmers, and higher-income households do not automatically reach those who need them most. The history of technology diffusion in developing countries is full of gaps between early adopters and the broader population that persisted for years or decades without active policy intervention.

Open-source AI tools, locally trained models, subsidized hardware for public health facilities, and community digital literacy programs can all widen access. But they require political will and sustained investment. Those are not technical problems. They are the actual constraints on equitable outcomes, and pretending otherwise does not help.

Figure 4 shows a side-by-side comparison of AI's estimated current and future effects on the economy and society in important areas, comparing the performance of developed economies with the current and 2030 target levels for LMICs.

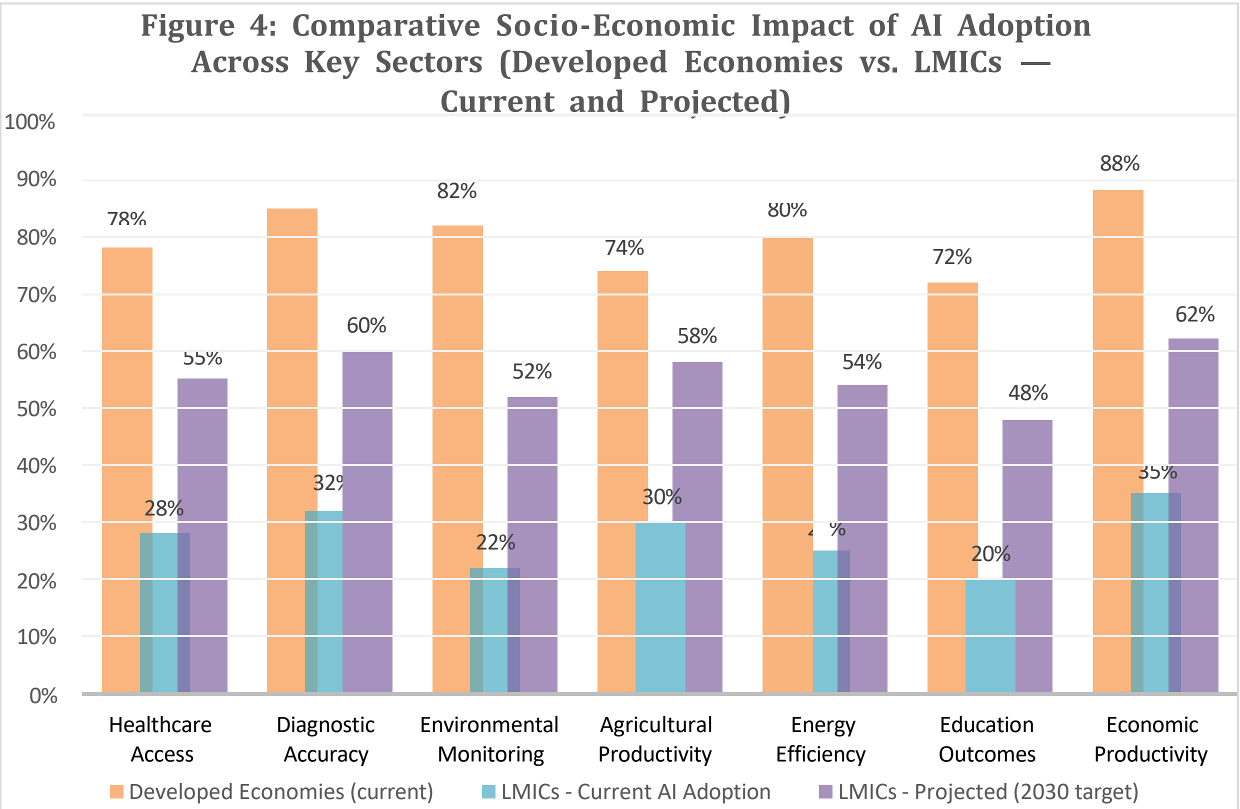

Figure 4: Comparative Socio-Economic Impact of AI Adoption Across Key Sectors — Developed Economies vs. LMICs (Current and 2030 Projection)

7. Challenges and ethical considerations

7.1 Data privacy and security

AI systems in healthcare and environmental monitoring collect and process data at scales that create novel privacy risks. Electronic health records, wearable biosensor streams, and GPS-referenced environmental readings are all sensitive personal information. Data breaches, unauthorized access, and secondary use without meaningful consent are not hypothetical risks. They have materialized in high-profile incidents involving both commercial platforms and public health systems.

For developing countries, the situation is compounded by weaker data protection legislation and limited enforcement capacity. Pakistan's Personal Data Protection Act (2023) established a legislative framework, but implementation remains a work in progress. Health data shared with AI platforms may sit on foreign servers, subject to foreign jurisdictions, and accessible to commercial actors under terms patients did not meaningfully consent to. Federated learning architectures, which enable model training without raw data leaving local servers, and differential privacy mechanisms can help. But they complement regulatory frameworks. They do not substitute for them.

7.2 Algorithmic bias and fairness

Machine learning models learn from historical data. Historical data reflects historical patterns of advantage and disadvantage. A predictive risk model trained on data from a health system where certain groups received less care will learn to predict lower risk for those groups. Not because they are at lower risk, but because their conditions were less frequently diagnosed and recorded. The model learns the bias in the records.

This is not an abstract concern. A widely cited study found that a commercial algorithm used to identify patients needing complex care systematically recommended less additional care for Black patients than for white patients with equivalent health needs. The effect traced back to using healthcare cost as a proxy for health need, which confounded race with differential access to care (Obermeyer et al., 2019). For applications deployed in developing countries, the training data is typically from high-income country populations with different demographics, disease patterns, and healthcare utilization. That matters.

Addressing algorithmic bias requires diverse training data, prospective bias auditing during development, ongoing performance monitoring across subgroups after deployment, and governance mechanisms that give affected communities a real voice in how AI systems are designed and evaluated.

7.3 Infrastructure and workforce constraints

The infrastructure requirements for AI deployment, reliable power, stable connectivity, adequate computing hardware, are not met in large parts of the developing world. Rural areas face the sharpest gaps, but urban areas in low- and middle-income countries also deal with outages, bandwidth limitations, and hardware procurement problems that constrain what is practically deployable.

The workforce gap may matter even more. Pakistan produces roughly 25,000 IT graduates per year, but skill distribution skews toward practical deployment rather than frontier AI research (PSEB, 2022). Building the capacity to develop, adapt, and maintain AI systems locally, rather than simply procuring foreign-built platforms and hoping they work, requires sustained coordination among universities, industry, and government. That coordination is harder to achieve than buying software, and it matters more.

Table 4 summarizes the main challenges covered in this chapter. For each one, it lists the challenge category and the mitigation strategies that current literature and policy evidence actually support.

Table 4: Key Challenges to AI Adoption in Developing Countries and Mitigation Strategies

Challenge	Category	Description	Proposed Mitigation
Algorithmic Bias	Ethical / Fairness	Training data underrepresents LMICs; performance	Diverse local datasets; mandatory subgroup performance reporting;

		gaps across demographic subgroups	community-inclusive design
Data Privacy & Security	Ethical / Legal	Sensitive health data exposed to commercial platforms; weak enforcement of data protection law	Federated learning; differential privacy; strengthen PDPA enforcement; data sovereignty clauses
Infrastructure Deficit	Technical	Power and connectivity gaps prevent deployment of real-time AI systems in rural areas	Edge AI; solar-powered sensor nodes; offline-capable mobile diagnostics
Lack of Interpretability	Technical / Trust	Black-box models reduce clinician and regulator confidence; accountability gaps	Explainable AI (XAI) methods; mandatory explanation requirements for high-stakes decisions
Workforce Shortage	Human Capital	Insufficient AI researchers and data scientists; clinical staff not trained to use AI tools	University AI programs; continuing professional education; public-sector AI fellowships
Affordability Barriers	Economic / Access	High device and platform costs exclude low-income populations and under-resourced health facilities	Subsidized device procurement; open-source model sharing; government-funded pilot programs
Regulatory Gaps	Governance	No AI-specific liability framework; unclear accountability when	Dedicated AI regulation; sandbox environments for pilot testing; international regulatory cooperation

		AI-assisted decisions cause harm	

Source: *Authors' synthesis. PDPA = Personal Data Protection Act (Pakistan, 2023); XAI = Explainable Artificial Intelligence.*

8. Future trends and opportunities

Foundation models, large pre-trained models adaptable to specific tasks with relatively small amounts of domain-specific data, may significantly reduce data requirements for new deployments. If a medical imaging foundation model pre-trained on large global datasets can be fine-tuned for local diagnostic tasks with a few thousand locally labeled examples, the economics of local AI development change substantially. Early evidence from models like Med-PaLM suggests this direction is viable (Singhal et al., 2023). It is early evidence, but it is real.

Edge AI, deploying trained models on local devices rather than cloud servers, addresses the connectivity and latency constraints that limit cloud-dependent applications. A diagnostic model running on a low-power device at a rural health post, without internet connectivity, changes what is possible where connectivity is unreliable. Advances in model compression, quantization, and hardware design are making increasingly capable models deployable on increasingly modest hardware. The trend line is encouraging.

Explainable AI (XAI) methods that make model reasoning interpretable to non-expert users matter both for clinical trust and regulatory compliance. Clinicians and policymakers are, reasonably, reluctant to act on AI recommendations they cannot interrogate. That reluctance is rational and should be treated as such, not as an obstacle to be managed. Progress in XAI may increase willingness to deploy AI where accountability demands transparency.

For Pakistan and similar contexts, the most consequential near-term opportunity is probably satellite and remote sensing data. Land use change, water body dynamics, crop health, urban heat island intensity, and air quality patterns can all be monitored at sub-national scale from satellite data. This does not require local ground infrastructure. Increasingly it is accessible through open platforms. That is a meaningful shift.

9. Conclusion

AI-enabled smart systems can genuinely improve healthcare outcomes, strengthen environmental monitoring, and support more evidence-based planning. The evidence for this, in specific applications and contexts, is credible. It would be dishonest to overstate it, and it would also be dishonest to dismiss it.

The record of technology adoption in developing countries justifies realistic expectations about pace, distributional patterns, and the conditions under which benefits actually materialize. For Pakistan in particular, the strategic priority is not passive adoption of systems designed elsewhere. It is active participation in shaping how those systems are designed, trained, and governed for local relevance.

That requires three things no algorithm can substitute for. First, data infrastructure that actually captures local populations and conditions. Second, human capacity to develop and critically evaluate AI tools, not just operate them. Third, governance frameworks that ensure accountability and protect the rights of those most affected by automated decision-making.

AI will not advance development goals on its own. It never was going to. Development still requires institutional quality, political commitment, and resource investment. What AI can do is make some of that investment more productive, some pressing problems more tractable, and some hard tradeoffs somewhat easier to navigate. That is enough to justify serious engagement, and enough to justify scrutiny of how that engagement unfolds.

References

Brownstein, J. S., Chu, S., Marathe, A., Marathe, M. V., Nguyen, A. T., Paolotti, D., Perra, N., et al. (2017). Combining participatory influenza surveillance with modeling and forecasting: Three alternative approaches. *JMIR Public Health and Surveillance, 3*(4), e83.

Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems, 29*(7), 1645–1660.

Javed, A. (2020). The scope of information and communication technology enabled services in promoting Pakistan economy. *Asian Journal of Economics, Finance and Management*, *2*(4), 1–9.

Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., et al. (2023). Learning skillful medium-range global weather forecasting. *Science, 382*(6677), 1416–1421.

Muralidharan, K., Singh, A., & Ganimian, A. J. (2019). Disrupting education? Experimental evidence on technology-aided instruction in India. *American Economic Review, 109*(4), 1426–1460.

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.

Oviedo, S., Vehí, J., Calm, R., & Armengol, J. (2017). A review of personalized blood glucose prediction strategies for T1DM patients. *International Journal for Numerical Methods in Biomedical Engineering, 33*(6), e2833.

Qureshi, H., Mahmood, H., Alam, S. E., Aabroo, A., Channa, Z. I., Shaheen, A., Fatima, Z. J., et al. (2026). Vaccination in delivery wards leads to high timely coverage of hepatitis B birth dose vaccine in Islamabad–Pakistan, 2023. *JPMA: The Journal of the Pakistan Medical Association, 76*(1), 86.

Rajpurkar, P., Chen, E., Banerjee, O., & Topol, E. J. (2022). AI in health and medicine. *Nature Medicine, 28*(1), 31–38.

Rodrik, D. (2018). *New technologies, global value chains, and developing economies* (NBER Working Paper No. 25164). National Bureau of Economic Research.

Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, *3*(5), 637–646.

Singhal, K., Azizi, S., Tu, T., Mahdavi, S. S., Wei, J., Chung, H. W., Scales, N., et al. (2023). Large language models encode clinical knowledge. *Nature, 620*(7972), 172–180.

Sjoding, M. W., Dickson, R. P., Iwashyna, T. J., Gay, S. E., & Valley, T. S. (2020). Racial bias in pulse oximetry measurement. *New England Journal of Medicine, 383*(25), 2477–2478.

Snyder, E. G., Watkins, T. H., Solomon, P. A., Thoma, E. D., Williams, R. W., Hagler, G. S. W., Shelow, D., Hindin, D. A., Kilaru, V. J., & Preuss, P. W. (2013). The changing paradigm of air pollution monitoring. *Environmental Science & Technology, 47*(20), 11369–11377.

Strohmeier, S. (2022). Artificial intelligence in human resources—An introduction. In S. Strohmeier (Ed.), *Handbook of research on artificial intelligence in human resource management* (pp. 1–22). Edward Elgar Publishing.

Topol, E. (2019). *Deep medicine: How artificial intelligence can make healthcare human again*. Hachette UK.

Turakhia, M. P., Desai, M., Hedlin, H., Rajmane, A., Talati, N., Ferris, T., Desai, S., et al. (2019). Rationale and design of a large-scale, app-based study to identify cardiac arrhythmias using a smartwatch: The Apple Heart Study. *American Heart Journal, 207*, 66–75.

Albu, D. (2023). World development report 2023: Migrants, refugees, and societies. *Drepturile Omului, 81*.