


Diagnosing

An abstract graphic featuring a network of blue lines and dots, with a series of overlapping, wavy orange lines in the center, creating a sense of depth and complexity.

Closing Gaps in
Healthcare Access with AI

Disparities



EDITOR-IN-CHIEF

ANGELA TUFVESSON

SUB-EDITOR

SUGANDHI RAVINDRANATHAN

PLATFORM DIRECTOR

METTE HALBORG THORNGAARD

ART DIRECTION & DESIGN

MAYA MORTENSEN-RAMLILL

SECRETARIAT

JULIE BUUR TRÆRUP

SEYI OLAWOYIN

CELINE HUMMELMOSE

THOMAS MARLAND

LENIA MIRA KESORE

MADE POSSIBLE WITH
SUPPORT FROM

AMGEN

FIND US AT

**NØRREBROGADE 45C
2200 COPENHAGEN
DENMARK**

INFO@RE-SOLVEGLOBALHEALTH.COM



From the editor

AI is rapidly transforming how we work, communicate, and make decisions, and healthcare is no exception. From faster diagnoses to extending care in resource-constrained settings, AI has vast potential to address some of the most complex challenges and barriers to healthcare access.

Yet it is difficult not to feel both fascinated by what these tools make possible and uneasy about what can be lost when judgment, context, and human experience are reduced to an algorithm. Or when AI tools are treated as a panacea for issues that are fundamentally human and systemic.

The arrival of AI in healthcare brings important questions about bias in datasets, infrastructure capability, data sovereignty, and environmental impact. An AI diagnostic may be highly accurate in one population but less reliable when trained on data that does not reflect the people it is intended to serve. AI clinical support systems could improve decision-making but be of little use in a facility without the infrastructure or workforce capacity to integrate it.

These are not reasons to resist AI, but impetus to get its foundations right to build more accessible, responsive, and equitable health systems.

ANGELA TUFVESSON
EDITOR-IN-CHIEF

Table of contents

Can AI become a trusted frontline diagnostician?

by Farai Shawn Matiashe

04

AI will augment, not replace, doctors

by George Zarkadakis

10

The push for local data to fuel Vietnam's medical AI

by Thu Dinh

12

AI raises the stakes for African health data protection

by Ray Mwareya

16

Why Africa should build trust in AI differently

by Caroline S Mbindyo

20

Data centre boom comes with steep health and climate costs

by Danny Buckland

22

How AI could crack vaccine science's hardest problem

by Roberto Amato & Bilal A Mateen

28

Better AI for women's health hinges on more data

by Becky McCall

30

Health AI must be designed with equity at its core

by Nyawira Gitahi

34

Is AI closing care gaps for underserved communities?

by Brian Mastroianni

38

Why health AI needs trust built into workflows

by Aryan Chaudhary

42

AI alone cannot fix Africa's health service delivery gaps

by Paul Adepoju

44

Can AI become a
trusted frontline
diagnostician?





By Farai Shawn Matiashe

In communities across Zimbabwe, AI-powered tools are helping frontline workers detect disease earlier and facilitate timely care. Scaling these technologies will depend on ensuring accuracy, building trust, and strengthening infrastructure.

When Janet Museta woke up in a hospital in Zimbabwe's Manicaland province, after being unconscious for hours, she was not as much worried about the extent of her injuries as about the risks of medical emergencies like stroke and high blood pressure that would impede her ability to work and take care of her family.

The 60-year-old mother of five had lost a leg after being hit by a speeding motorcycle at a shopping centre in her home village of Gwanyamwanya outside Nyanga, a popular tourist town north of Zimbabwe's eastern highlands.

After surgery in February, Museta learned how to use her new artificial limb. But after a month, her health took a turn for the worse. She started feeling dizzy and experiencing severe chest pain.

A digital health promoter—a community health worker who uses digital tools to provide basic screenings and health education—came to her house to measure her blood pressure, revealing that it was extremely high. Museta spoke to a doctor through a video call on a smartphone, who diagnosed her with hypertension and prescribed medication, which Museta takes once a day.

"I believe my blood pressure was high because of stress. It has not been easy," says Museta, who survives on farming. "Imagine waking up one day, realising that you cannot do household chores or go to the farm."

Alarmingly, she then started experiencing pain on the right side of her abdomen and numbness—warning signs of an imminent stroke. The digital health promoter used a digital tool called Afya Selfie to assess her stroke risk. Thankfully, it was found to be minimal. "I was relieved. I feared the worst," Museta says.

AI smartphone screening tool Afya Selfie assesses Janet Museta's risk of having a stroke near Nyanga, Zimbabwe.
Photo: Farai Shawn Matiashe

Afya Selfie is a Zimbabwean AI screening tool that uses a smartphone camera and photoplethysmography—a technique that measures blood volume changes using light—to estimate indicators such as blood pressure, heart rate, and stress levels. Through a selfie scan, the app assesses cardiovascular risk. Patients pay between US\$1–3 for screening.

It is one of a growing number of AI-powered technologies that are expanding digital access to health assessments across Zimbabwe and beyond, bringing speedy diagnosis and specialist support to underserved communities.

AI-enabled diagnostic tools

In March, Zimbabwe launched its National Artificial Intelligence Strategy (2026-30), aimed at accelerating the adoption of AI in different sectors, including health. Against this backdrop, innovators have been deploying a growing number of AI tools such as Afya Selfie to serve as diagnostic aids, particularly in rural areas.

Dr Admore Jokwiro, a co-founder and chief medical officer at ZimSmart Villages, which developed Afya Selfie, says the AI tool works by combining photoplethysmography, smartphone-camera imaging, machine-learning models trained on large clinical datasets, and conversational AI interfaces to assess a patient's risk profile in real time.

"Together, these tools function as a...first line of clinical intelligence that can triage, flag, and refer patients who would otherwise go undetected for years," he says.

Zimbabweans have been grappling with an underfunded health system for decades. There are shortages of not only basic medicine like paracetamol but also of health workers, including doctors, specialists, and nurses. The ratio of physicians to the population is critically low; specialist services are concentrated in urban centres, and rural and semi-urban communities face a rising burden of non-communicable diseases like cardiovascular disease, diabetes, hypertension, anaemia, and different types of cancer.

"AI tools address that gap directly by bringing screening to the community rather than waiting for the community to reach a clinic," Jokwiro explains.

To date, nurses and digital health promoters have used Afya Selfie's cardiovascular risk predictor to screen more than 5,000 people, with a striking 30–35% returning abnormal results like high hypertension or cardiovascular disease risk. The results show that

for every three people screened, at least one is carrying a metabolic or cardiovascular risk factor that was previously invisible to the health system, says Jokwiro.

ZimSmart Villages has since expanded its suite of AI-powered screening tools. Its breast cancer detection app uses image analysis algorithms to identify suspicious lesions from photographs taken with a handheld ultrasound scan. Using the technology, experts have assessed more than 3,000 women, detecting suspicious findings in around 10% of cases and referring them for further investigation.

Another tool, Anemia Checker, estimates haemoglobin levels non-invasively by analysing images of the lower part of the eye with a smartphone, offering an alternative to laboratory-based blood testing. Project Tendesai, meanwhile, screens for depression, trauma, and substance use disorders using validated psychometric assessments delivered in local languages.

Speed and accuracy of AI tools

Professor Tawanda Mushiri, executive director at the Scientific & Industrial Research & Development Centre, which provides technological solutions for sustainable development in Zimbabwe, says that although the accuracy and reliability of AI in diagnostic tools vary enormously, there are cases where they have outperformed traditional methods of providing healthcare.

One example, he says, is the non-invasive anaemia detection application (NiADA) developed by Dawa Health, a digital health provider operating in Zimbabwe, Rwanda, and Zambia. The AI-powered tool estimates haemoglobin levels without requiring a blood test, enabling healthcare workers to identify anaemia in pregnant women before childbirth and helping prevent potentially life-threatening complications. "Dawa Health's tools reportedly cut anaemia detection from three days to five minutes in rural pilots," Mushiri says.

Jokwiro says his organisation's screening tools have been validated against established clinical reference standards and have demonstrated strong performance. He says the tools correctly identify people with a health condition and accurately rule out those without one in more than 85% of cases. According to him, this level of accuracy is comparable to the clinical standards routinely achieved by diagnostic tools used in well-equipped hospitals in high-income countries.

"Sensitivity above 85% means the tool correctly identifies the vast majority of individuals who truly



An AI-powered application was used to monitor Annamore Murenzvi's pregnancy for her twins in Huana, Zimbabwe.
Photo: Farai Shawn Matiashe

have the condition being screened for, minimising the risk of false reassurance," Jokwiro points out.

Building trust in AI as a healthcare assistant

Securing the trust of health workers and patients in AI's capabilities is often just as important as developing the technology itself. When Museta was told that the risk of having a stroke in the future was minimal, she was sceptical and struggled to believe a smartphone could provide such a reliable assessment. "I asked the health worker a lot of questions and her answers were convincing," Museta says.

Annamore Murenzvi, from Hauna, a farming community next to Nyanga, came for a pregnancy ultrasound scan in October last year. Using a tablet's camera, an AI-powered tele-obstetrics app developed by ZimSmart Villages detected twins and that she was already in her second trimester. Health workers monitored her until she gave birth to two

bouncing baby boys early this year. "I did not have any complications. I have always had trust issues with scans, but this one was accurate," says the first-time mother.

"I have confidence in this AI tool. It helps us check the health of both the mother and the foetus. It reduces anxiety for the mother."

Using the app, nurses perform obstetric scans using portable devices, supported by remote specialists. The app tracks gestational progress, predicts preeclampsia risk, and monitors foetal wellbeing.

Memory Difara, a registered nurse and midwife at ZimSmart Villages' telehealth centre, says that when Murenzvi came for the scan, the app did not detect any abnormalities.

"The foetal heart rate was normal. We screened, using the app, for major birth defects like anencephaly, cleft lip, and spina bifida, which were normal. The placenta was healthy, umbilical cord and amniotic fluid were normal," she says. "The AI app

calculated the gestational age and estimated foetal weight and date of delivery. This helped the mother to plan for the delivery of her babies.”

Difara believes the diagnostic tool is accurate. “I have confidence in this AI tool. It helps us check the health of both the mother and the foetus. It reduces anxiety for the mother,” she says.

Jokwiro emphasises that the AI tools are designed as screening instruments, not definitive diagnostic replacements, explaining that their role is to identify people who need further investigation, rather than replace a clinician’s final judgment.

Mushiri agrees that the tools work best as decision support, not replacement. “The framing that keeps recurring in the literature is AI as an assistant to a health worker, not a substitute for one. AI-driven decision-support systems can act as intelligent assistants that enhance clinical reasoning, promote adherence to guidelines, and improve accuracy, rather than [replace] healthcare professionals,” he says.

Towards a long-term vision

Despite their potential, integrating these AI tools into Zimbabwe’s health system at scale remains challenging. Jokwiro notes that connectivity and device access remain a tough call. He says while smartphone penetration is growing rapidly, the most remote and vulnerable communities still lack reliable internet access and affordable devices.

Mushiri says regulatory and accountability gaps persist. “It is unclear how the Medical and Dental Practitioners Council of Zimbabwe can regulate an AI-based app that diagnoses and issues prescriptions,” he says, noting there is no clear regulatory and liability framework.

A genuine AI health infrastructure for diagnostics in Zimbabwe and beyond will rely on open standards, locally owned data infrastructure, and a training pipeline.

Mushiri adds that while trust is a known adoption barrier, workforce training gaps make it worse. “Healthcare workers often lack training in interpreting algorithmic outputs or integrating machine-learning insights into clinical practice,” he says, explaining that health workers operating a genuinely accurate tool can misuse it simply because no one trained them on what the output means.

Jokwiro, as a health data expert, believes that a genuine AI health infrastructure for diagnostics in Zimbabwe and beyond will rely on open standards, locally owned data infrastructure, and a training pipeline. This pipeline will continuously improve performance by using African patient data.

Mushiri says a practical starting point is acknowledging that infrastructure, not algorithms, is the bottleneck. He highlights that there is a need to enhance devices and connectivity first and implement low-data, offline, and localised models.

To keep costs down, tools such as Afya Selfie and the tele-obstetrics app are deployed through existing infrastructure, with ZimSmart Villages transforming 12 state-owned post offices into telehealth centres across the country. Nurses based at these centres use AI-powered tools to screen patients, while people who need further care can consult specialists remotely through telehealth services. “In the immediate term, the priority is scaling the post office clinic model,” Jokwiro says, noting that there are more than 200 post offices across Zimbabwe that could be used.

He says the longer-term vision is a national AI health screening network that feeds anonymised, aggregated, epidemiological data into the Ministry of Health’s planning systems, giving policymakers, for the first time, a real-time picture of disease burden at the community level rather than relying on hospital admission data, which captures only the most severe cases.

For patients like Museta, the promise of this technology is not just better data for policymakers, but earlier detection and reassurance at a time when uncertainty can be frightening. “I thought I was dying, not of the injuries but of a stroke because of stress. I could not sleep. But I am feeling better now,” she says.

A digital health promoter using a device to measure Janet Museta’s blood pressure near Nyanga, Zimbabwe.
Photo: Farai Shawn Matiashe



AI will augment, not replace, doctors

The opinions expressed are those of the author and do not necessarily reflect the position of Re:solve Global Health.

Generative AI is moving healthcare beyond isolated algorithms towards general medical intelligence. The challenge now is not whether medicine will adopt AI, but how to ensure it does so safely, equitably, and for the benefit of patients.

AI has been transforming medicine for years, although most people only became aware of its potential after the arrival of generative AI in the form of large language models (LLMs), such as ChatGPT and Claude. Long before generative AI entered public consciousness, machine learning was already helping radiologists detect cancers, supporting clinicians in interpreting medical images, and accelerating drug discovery. Yet these successes were largely confined to individual tasks. Each AI system was designed to perform one specific function exceptionally well but had little ability to transfer its knowledge elsewhere. That era is ending.

Generative AI represents a fundamental change in the relationship between humans and machines. Instead of creating specialised tools for isolated clinical problems, we can build systems capable of understanding language, integrating diverse medical information, and assisting physicians across an entire clinical workflow. The implications extend far beyond automation. They point towards a new cognitive infrastructure for healthcare.

The need for AI has never been greater

Healthcare systems everywhere face the same convergence of pressures: ageing populations, chronic workforce shortages, escalating costs, and an explosion of medical knowledge. At the same time, clinicians are expected to synthesise information from electronic health records, imaging, laboratory tests, genomics, wearable devices, and an ever-growing body of scientific literature.

No individual physician—regardless of experience or expertise—can realistically process all this information unaided. In my discussions with physicians, one of the most consistent complaints is that they are overburdened with administrative tasks that prevent them from spending more time with their patients.

This is precisely where AI creates value. Rather than replacing clinicians, AI functions as a cognitive multiplier. It helps organise complexity, take over repetitive low-value tasks, and identify relevant patterns and surface evidence that would otherwise remain buried beneath overwhelming volumes of data. This distinction matters. The most successful applications of AI have not attempted to remove doctors from the decision-making process. Instead, they improve physicians' ability to make informed decisions more efficiently.

Radiology offers perhaps the clearest example. Modern deep learning systems have demonstrated remarkable performance in detecting disease while simultaneously reducing clinician workload. Studies increasingly show that the combination of human expertise and AI assistance consistently outperforms either working alone. Generative AI expands this partnership even further. Large language models are uniquely suited to medicine because healthcare is fundamentally organised around language. Patient histories, clinical notes, referral letters, research papers, treatment guidelines, and multidisciplinary discussions are all expressed through language.

For the first time, AI systems can read, interpret, summarise, and reason across these diverse sources of knowledge while communicating with clinicians in natural language. This changes the economics of expertise. Instead of requiring every physician to memorise an ever-expanding medical literature, clinicians gain immediate access to continuously updated knowledge. Instead of searching through hundreds of publications, they can rapidly obtain evidence summaries tailored to the clinical question before them.

In resource-constrained health systems, the implications could be transformative. Countries facing shortages of specialists may eventually provide frontline clinicians with AI systems capable of offering expert-level decision support. Such tools could help reduce geographical inequalities in healthcare by making specialist knowledge available wherever patients are treated.

But we should resist the temptation to frame this as a race towards replacing physicians. Medicine is not

simply pattern recognition. Clinical judgement depends on empathy, ethical reasoning, contextual understanding and trust—qualities that remain fundamentally human. Patients do not merely seek accurate diagnoses; they seek reassurance, shared decision-making, and compassionate care. AI should therefore augment these uniquely human capabilities rather than compete with them. Accordingly, the next frontier lies in what researchers increasingly describe as generalist medical AI.

Sound governance underpins equitable access

Unlike today's AI models, future systems will integrate multiple forms of medical information simultaneously: clinical text, diagnostic imaging, laboratory data, genomics, multiomics, and physiological signals from wearable devices. They will reason across these modalities while continuously incorporating new medical knowledge. Such systems could support clinicians throughout diagnosis, treatment planning, chronic disease management, and preventive care. Yet their success will depend less on technological capability than on responsible governance.



In radiology, modern deep-learning systems show improved performance in detecting disease while simultaneously reducing clinician workload.
Photo: Tima Miroshnichenko

The greatest risks are not science-fiction scenarios involving autonomous machines. They are far more immediate: biased training data, opaque algorithms, unequal access to computational infrastructure, inadequate validation, and insufficient regulatory oversight. Medical AI must earn trust in the same way that medicines and medical devices do: through rigorous evaluation, transparency, and continuous monitoring.

Equally important is ensuring equitable access. Training advanced AI systems requires enormous computational resources that remain concentrated among a handful of technology companies. Clinical and genomic data used to train AI systems are predominantly sourced from advanced Western societies and are heavily biased towards specific human phenotypes and genotypes, thus excluding and marginalising much of the world's population. Without deliberate efforts to support open science, access to more diverse and shared medical datasets, and international collaboration, AI risks widening global health inequalities rather than narrowing them.

Healthcare has faced technological revolutions before—from antibiotics to imaging and genomics. Each transformed medicine not because it eliminated physicians, but because it expanded what physicians could achieve. Generative AI should be viewed in exactly the same way.

The future of healthcare is unlikely to belong to doctors alone, nor to AI acting independently. It will belong to clinicians working alongside intelligent systems that enhance their expertise while preserving the human relationship at the heart of medicine.

The real revolution is not AI replacing doctors. It is AI helping doctors become even better at what only humans can ultimately do: care for patients. ■

Re:

George Zarkadakis has a PhD in AI in medicine and is chief innovation officer at Xeomics and an author and futurist specialising in AI, digital health, and emerging technologies. His work focuses on the responsible adoption of AI in healthcare, helping organisations translate technological innovation into improved patient outcomes while addressing the ethical, regulatory, and societal challenges of medical AI.

The push for local data to fuel Vietnam's medical AI



Vietnamese medical developers are reducing their reliance on datasets drawn from Western populations to build AI tools grounded in local realities. This effort will enable new digital health systems to be accurately tailored to the people they serve.

Something was not right with two-year-old Binh's* medication. The little boy from Hanoi was battling Kawasaki disease, a rare condition that can dangerously widen the blood vessels around his heart. To keep him safe from complications, doctors prescribed him clopidogrel, a medicine designed to prevent blood clots.

For two months, the toddler was given the drug as prescribed. But when his parents brought him back for

Local patient data is helping Vietnamese developers build AI predictive models that help deliver personalised care to children with rare diseases.

Photo: Beth MacDonald

follow-up checks, his blood-test results had still not stabilised. The medication, a standard treatment used successfully by millions around the world, appeared to be ineffective for him.

The answer lay in his DNA. A genetic test conducted by GeneStory, a local medical startup, along with the doctor's assessment, revealed that Binh's unique genotype made it impossible for his body to metabolise clopidogrel properly. His condition finally stabilised when his doctors switched him to an alternative treatment.

For Dr Sy Nam Vo, co-founder and chief scientific officer at GeneStory, the toddler's experience is a reminder that a one-size-fits-all approach to medicine does not always work.

The startup's AI predictive models were made possible by a landmark project, conducted by Vo and his colleagues at a research institute, to sequence 1,000 Vietnamese genomes. Without this foundational dataset, developing an AI application capable of catching such biological mismatches would have been impossible.

"If a medical AI is trained exclusively on foreign datasets such as Western or Caucasian populations, it might blindly recommend standard treatments that might be ineffective, or actively harmful, to Vietnamese patients," Vo explains.

To safely treat patients like Binh, Vietnamese developers and hospitals are moving away from Western datasets, seeking ways to responsibly collect and utilise local patient data.

The AI wave in Vietnamese healthcare

Vietnam's medical AI boom is unfolding within a broader, government-led effort to digitalise healthcare and move towards universal health coverage. The National Digital Transformation Programme, approved in 2020, identified healthcare as one of the priority sectors. A year later, the National AI Strategy set goals for developing and using AI by 2030. In healthcare, this includes improving the sharing of health data and using AI to support diagnosis, treatment decisions, remote monitoring, personalised care, and drug research.

Against this backdrop, Dr Pham Duc Phuc, a public health expert at the Institute of Environmental Health and Sustainable Development, describes the current scene as "a hundred flowers blooming", with tech firms and hospitals racing to create new medical AI applications.

They are focusing on diagnostic imaging, predicting treatment outcomes, community screening, virtual medical assistants, health record analysis, and support for drug and therapy research.

"In the context of our current doctor shortage and the massive amount of visits clinicians handle every year, AI is essential. A doctor supported by AI will work much more efficiently and deliver better care."

"Medical AI acts as a much-needed lifeline for an overwhelmed workforce," Tien Zung Nguyen, a professor at the University of Toulouse, France, adds from his experience testing AI-aided disease detection tools developed by his startup Torus AI in Vietnamese hospitals.

Vietnam notched up approximately 15 doctors per 10,000 people in 2025, up from just 9.81 in 2020, the Ministry of Health reported at a conference at the end of 2025. However, the ratio of doctors still falls short of actual demand. Major centres like Hanoi's Bach Mai Hospital report operating at up to 170% of their original design capacity, placing doctors under constant pressure as they cope with heavy patient loads.

"In the context of our current doctor shortage and the massive amount of visits clinicians handle every year, AI is essential," Tien Zung Nguyen explains. "A doctor supported by AI will work much more efficiently and deliver better care."

Hidden risk of foreign data

Building reliable medical AI demands one essential ingredient: local patient data. In Vietnam, however, this information lies in a fragmented, disconnected web.

"These foreign resources could miss the unique genetic makeup of Vietnamese patients, holding back research into personalised care."

Vietnam's healthcare system serves a population of 102.3 million through a vast network extending from central referral hospitals to local primary care. A 2020 World Bank report describes a fragmented digital landscape in which hospital records, health insurance claims, vaccination registries, and communicable disease databases are stored on separate platforms built to different standards. Hospitals commonly operate their own record systems, while commune facilities may use several different programs; records can even be duplicated within the same institution. At one point, Vietnam was dealing with around 1,000 different software systems just for laboratory results.

This fragmentation makes it difficult to compile the large, high-quality datasets needed to train AI models. In genetics, a lack of local data forced early researchers to rely heavily on international databases built mostly from people of European descent.

"These foreign resources could miss the unique genetic makeup of Vietnamese patients, holding back research into personalised care," Vo points out. Bridging the gap, his team had to build a research system from the ground up, figuring out everything from processing raw samples to analysing complex data.

When medical AI learns mostly from Western patients, it learns to spot diseases exactly as they look in those specific groups. Drop that same system



Tech firms in Vietnam are building AI models for health specifically tailored to the Vietnamese population.
Photo: GeneStory

into a Vietnamese hospital, and it can easily misread deep differences in genetics, local diets, and common illnesses.

Obesity conundrum

Pham cites obesity as a primary example. Consider a rural clinic where a physician examines a patient who appears slim and scores within a healthy range on standard Western body mass index (BMI) charts. These imported metrics, however, are often ill-suited to local demographics. Many Vietnamese patients carry visceral fat around their internal organs, a condition Pham describes as "thin on the outside, fat on the inside". An AI tool trained exclusively on Western datasets would likely misidentify these individuals as healthy, potentially overlooking critical early warning signs of metabolic disease.

This same problem is happening across Southeast Asia. A regional study revealed that neighbouring countries are battling the same messy records and growing fears of AI bias. One Thai policymaker, for instance, noted that several AI tools failed noticeably when applied to local patients, despite working flawlessly on their native Western datasets.

Scramble for datasets

To break this reliance on foreign data, Vietnamese tech teams are scrambling to build homegrown datasets. However, getting that raw data is not easy.

While existing laws permit the use of health data for research, Pham notes that a clear, standardised pathway for accessing hospital archives remains elusive. Medical imaging serves as the most accessible data source for AI developers in Vietnam, as these files follow standardised formats and are easily anonymised.

Conversely, researchers face significant barriers when attempting to access other types of health records. "Hospital directors worry about strict privacy regulations and navigating multiple layers of approval," Pham notes. "Fearing potential data leaks, they often choose the safest path: locking their files away entirely."

"I believe AI is the future, and it will be part of every aspect of our lives. Using it is unavoidable. I just think they [the developers of medical AI models] should keep user data protection in mind while developing their products."

Public trust is another delicate issue. Hanh Nguyen, a 28-year-old office worker who frequents a private clinic in Hanoi, says she would not feel comfortable sharing her medical history unless she was completely sure her detailed information would be kept safe.

"I believe AI is the future, and it will be part of every aspect of our lives. Using it is unavoidable. I just think they [the developers of medical AI models] should keep user data protection in mind while developing their products," she says.

Consequently, a national data drought persists. The fragmented, homegrown datasets currently available remain confined within isolated partnerships. Pham warns that this gap risks fostering new forms of medical bias within Vietnam. If AI models are trained exclusively on data from a few central, urban hospitals, they will reflect only a narrow demographic, failing to capture the diversity of the Vietnamese population across different ages, geographies, and physical characteristics.

A homegrown data foundation

Turning Vietnam's medical AI dreams into reality will require a major team effort, backed by strong government support, to finally piece together the country's scattered health records.

Partnerships are emerging, with hospitals, universities, and technology groups sharing data for specific projects. Yet, as Tien Zung Nguyen argues, this teamwork must now become systematic.

Research teams specialising in the same field, such as cardiology, should pool their data and develop AI tools together rather than building separate systems in isolation.

Pointing to his own team's ongoing work with local doctors, Tien Zung Nguyen notes that their current pool of Vietnamese data remains quite small. "But as we push forward, it will steadily grow," he adds.

Indeed, experts are calling for a united effort to build a shared, national dataset. This foundation should truly reflect the Vietnamese people, weaving together the diverse traits of city, rural, and highland communities, while proportionally representing the country's 54 ethnic groups.

The shift begins on the ground, where hospitals are slowly trading their towering stacks of paper files for electronic health records. The government is driving a massive campaign to digitalise everyday healthcare, rolling out electronic health books and linking patient data directly to a national digital ID app. A patient's medical history will be securely carried on their smartphone, ready to be accessed by any doctor, anywhere.

This digital shift feeds into a much larger ambition. The Ministry of Health is laying the groundwork by building a National Health Database, a massive effort designed to smoothly connect these new digital records across the entire nation.

The ministry launched an intense 90-day nationwide campaign running from July to September. The mission is to meticulously scrub, standardise, and secure 12 major specialised healthcare databases. The initiative also locks down information security to ensure this sensitive data can be safely shared across the national network.

Pham advocates for the creation of a dynamic, coordinated data ecosystem—a unified national

network anchored in shared standards, collaborative frameworks, and rigorous ethical guidelines. "Success depends on hospitals fostering mutual trust and policymakers establishing clear, robust privacy regulations," he explains.

A growing number of local institutions are releasing their cleaned-up datasets to the wider research community. One example is a recent publication in *Nature Communications* about the 1,000 Vietnamese Genomes Project by Vo and his colleagues. Born from a private sector effort, it provides a much more accurate genetic baseline for the local population.

Vo acknowledges the current sample remains small compared to international projects. Expanding it will require long-term investment, broader participation, and stronger state support.

Besides building a broad data network, Tien Zung Nguyen believes Vietnam must develop its own language models and specialised medical AI to independently control the technologies and information on which its health system increasingly relies.

"Building domestic AI capability must be viewed not just as a technical upgrade, but as a matter of digital sovereignty," he says.

Above all, it is hoped these shared efforts will ultimately bring comfort and better care to patients themselves. Knowing that new diagnostic tools will be built on data that truly reflects her own people, office worker Hanh Nguyen feels a sense of hope about the future of domestic health technology.

"I'm not afraid to try them," she says, "as long as they remain under the careful control of doctors."

*Name changed to protect the patient's identity



The digital shift in healthcare starts with integrating electronic health books to build a shared national dataset.
Photo: GeneStory

AI raises the stakes for African health data protection



Paper data repositories present a great challenge for pooling health data and tracking a patient's medical history; digitising public health data is an effective solution. Photo: RDNE Stock Project

Health data is becoming a strategic resource in the AI era. In South Africa and across the continent, the question is not simply who has access to data, but who controls it, who benefits from the technologies built from it, and whether countries generating the data have the capacity to use it themselves.

In the subterranean room of a public clinic in Diepsloot, one of the most populous and lowest-income townships in Johannesburg, South Africa, lie boxes stacked with years of patient histories, all damp, greying, and decaying.

"This is a data repository—how we have always kept paper records over the last 100 years in South Africa's public clinics—torn manilas, water-rusted booklets,

clinical notes [becoming] unreadable over time," explains Bhaizani Bonza, a veteran nurse at Diepsloot South Clinic.

Paper health databanks like this are common across most public hospitals that serve low-income communities in South Africa, a nation of 60 million. If a woman is treated for maternal syphilis in Gauteng, the country's highest-income province, and relocates to low-income Eastern Cape, doctors would struggle to draw the pattern of her medical history, Bonza says.

From 2024, a nascent revolution has been underway, as a trial model, to discard South Africa's hospital data's "paper prison" dilemma and replace it with an integrated digital architecture known as the Electronic Medical Record (EMR) system, or e-Impilo (a Zulu word meaning life and health).

This is a blueprint for digitising public health data that the rest of Africa has also been working towards, with

varying degrees of success. But as health data becomes an increasingly valuable resource in the age of AI, a pressing question is emerging: who gets to decide how that data is used?

The America First Global Health Strategy, launched by the Trump administration in 2025, has prompted debate over US efforts to expand access to African digital health information, including electronic medical records, epidemiological data, and high-risk pathogen samples.

“Africa has high-quality raw data—hence the US request for Africa’s medical pathogens and genetic data mirrors the deeply entrenched patterns of old extractive industries in metals, forests, and ocean products.”

Countries like South Africa that have not signed a memorandum of understanding risk losing access to US-backed health programmes and funding streams. Those that do sign must weigh the benefits of continued health cooperation against questions about data governance, sovereignty, and who ultimately controls and benefits from the information generated by their populations.

Risk to African data sovereignty

The America First Global Health Strategy places disease surveillance and health data at the centre of US global health cooperation. The strategy has produced bilateral agreements with more than 30 partner countries so far, most of them low- and middle-income and a majority in Africa. These agreements aim to strengthen health security by improving access to health information, but there are significant concerns about data ownership and whether countries generating valuable health data will retain control over how it is used.

The legal texts underpinning these bilateral requests also do not make clear whether African countries will receive reciprocal access to US medical, pathological, and genetic data to train their own AI models and support homegrown biopharmaceutical innovation.

“It’s a classic divide-and-rule strategy, the most potent weapon in the colonial playbook,” Tatenda Chatikobo, a research fellow at the University of Warwick’s School of Law, says of the asymmetrical powerplay that has resulted in countries like Zimbabwe refusing a US\$367 million offer and Ghana a deal worth US\$109 million due to worries about digital sovereignty and data protection. Judges in Kenya briefly blocked the country’s US\$2.5bn agreement due to lawsuits over medical privacy.

Chatikobo is part of a Wellcome-funded project focusing on global health law, bioethics, epidemiology, and governance and co-authored the widely cited 2024 journal article, “Colonialism in the new digital health agenda”, in *BMJ Global Health*. The proposed deals are “extractive,” he says. “It’s not really sharing because the corresponding value is not really [clear]. We understand there’s vague talk of an exchange of aid on top of data.”

Valuable source of medical data

Africa has become an increasingly valuable source of medical data because many of the datasets used to develop AI tools today are heavily biased towards Western populations. They often fail to capture Africa’s genetic diversity, making African data essential for building more accurate and reliable AI tools.

“Africa has high-quality raw data—hence the US request for Africa’s medical pathogens and genetic data mirrors the deeply entrenched patterns of old extractive industries in metals, forests, and ocean products,” Chatikobo says.

Dr Joe Phaahla, the South African Deputy Minister of Health, speaking in his personal capacity, says Africa is home to some of the world’s highest genetic diversity. He believes the continent is sitting on a “treasure trove” of genomic and medical data that is “highly precious”. Domestic governments have been





A US-Africa collaboration should empower African clinical scientists, enrich local clinical skills, and keep a vast amount of intellectual property in Africa.
Photo: National Cancer Institute

collecting health information in various forms for decades, he says, and are now increasingly linking those records to digital biometric identifiers such as fingerprints and facial scans. “Our data could really unlock a lot of things,” Phaahla adds.

The Trump administration has also proposed providing African health agencies with proprietary, US-hosted digital health software at no cost. Critics, however, argue that the licensing arrangements deserve closer scrutiny. If health data must be stored on cloud servers outside Africa and governed by foreign legal jurisdictions, countries could have less practical control over how that data is accessed, managed, and reused.

Precision medicine and commercial collaborations

Donrich Thaladar, a prominent South African specialist litigator in reproductive law cases, is currently principal investigator of a National Public Health Institute of South Africa-funded project that examines the legal aspects of data science in health innovation in Africa. He takes a contrarian view to wholesale criticism of the US requests.

Thaladar says Africa has long argued that precision medicine has been shaped by research and innovation that underrepresent people of African ethnicity, leaving many populations without treatments tailored to their genetic diversity. Now, as

African countries build valuable genomic datasets, he believes the continent needs to explore methods to share this information with global biomedical innovators while ensuring that the resulting discoveries benefit African populations.

“If African populations are to have access to precision medicine, then they must offer their data so that research can be done on their data,” Thaladar says, adding that he believes the continent must move beyond the traditional aid-based model towards more equitable commercial collaborations. “For example, tell the US: ‘You can have access to this [information] on the following conditions—licence payment and any [intellectual property] generated from it, [to which] we must have certain access.’”

Thaladar says countries have an ethical responsibility to ensure high-quality data about their populations exists and that this data is shared with researchers who can use it to advance knowledge and improve health outcomes. “From a public health perspective, we want our populations to be healthy and have access to the best possible, newest kind of treatments and medicines,” he says.

For Thaladar, part of setting ethical and legal benchmarks for this kind of collaboration means insisting that African clinical scientists do more than draw blood for genetic samples. They should be empowered to extract the DNA, sequence the genetic data and map it, and store the data locally.

That kind of model creates meaningful jobs on the continent, enriches local clinical skills, and keeps a lot of the intellectual property at home in Africa. “Hopefully from there, a local bioeconomy can flourish,” he argues.

Plan to protect health data

Between fears of extraction and calls to negotiate, what is clear is that Africa must have a plan for its health data. What is needed, Deputy Minister of Health Phaahla argues, is for African countries to have a clear strategy for protecting and using their health data. “To refuse the Americans, or grant them access, we need to first ask ourselves: what actually are we going to do with this data as Africans to safeguard our interests?”


That question is arguably sharpest for countries without a digitised health data system. The US could offer these countries a lucrative sum to build a digital system from scratch, then attach a clause ensuring that once the repository is live, the US government and allied organisations get priority access to the genetic data it holds.

Photo: Towfiqu Barbhuiya

South Africa’s new EMR system offers a contrast as the country is pursuing its development aggressively and largely on its own terms. Indeed, South Africa is, arguably, the best-positioned African nation to push back against the America First Global Health Strategy. Its Protection of Personal Information Act (POPIA)—a robust framework for protecting citizens’ personal data, similar to the EU’s General Data Protection Regulation—gives government and courts legal grounds to resist a wholesale handover of health data to the US.

“We would need to give our nod of approval in a sort of referendum as citizens before our government can do a wholesale handover of sensitive health data to Americans,” explains Bonza, the nurse at Diepsloot South Clinic. “Such preparedness in terms of domestic data protection is arguably one of the reasons why the US has not had an easy time trying to arm-twist us in South Africa.”

For the rest of Africa, the current US strategy marks a sobering moment and a time for reflection, Chatikobo says. “We are yet to come up with very much African-owned, AI-driven health data tech infrastructure—and that’s a weak link.”



“We would need to give our nod of approval in a sort of referendum as citizens before our government can do a wholesale handover of sensitive health data to Americans.”

Why Africa should build trust in AI differently

The opinions expressed are those of the author and do not necessarily reflect the position of Re:solve Global Health.

Trust in AI depends on more than technology. Africa's opportunity is not simply to deploy AI diagnostics, but to show how trustworthy AI is built: through strong institutions, local evidence, and continuous evaluation.

In eight states across Nigeria, mobile teams have spent the past several years bringing ultra-portable digital X-ray units to hard-to-reach communities, screening tens of thousands of people for TB using CAD4TB, a computer-aided detection algorithm that flags likely disease on a chest radiograph in seconds. It is, on paper, exactly the kind of tool global health has long hoped for: a way to put specialist-level triage in front of patients who may never see a radiologist.

Independent validation studies in Tanzania and South Africa have shown that CAD4TB can accurately distinguish TB from healthy lungs, and its use aligns with the World Health Organization (WHO) 2020 recommendation that computer-aided detection can serve as an alternative to human interpretation for TB screening.

Yet a 2024 comparison of commercially available chest X-ray algorithms on the same South African population found CAD4TB performing meaningfully below several newer competing products, including Lunit and Nexus. CAD4TB was not fraudulent; nor was it unsafe. It remained a clinically useful tool. What changed was the availability of better comparative evidence.

This is the real state of AI diagnostics in Africa: not a story of broken algorithms, but of incomplete information. A health ministry procuring a TB screening tool, or a clinician relying on one, typically sees a single vendor's validation study, not an independent, head-to-head comparison showing how multiple products perform on local populations as technologies evolve. This is an information and governance gap, not simply an algorithmic one. And it reframes the trust question.

Much of the global conversation about AI diagnostics still asks whether we can trust AI. Africa should ask a different question: what kind of health system produces trustworthy AI? Trust is not something that can be built into an algorithm. It is earned through institutions that continuously evaluate, govern, and improve how AI performs in real-world care.

Clinical usefulness is a systems outcome

Health technologies rarely transform health systems on their own; systems determine the extent to which technology's promise becomes realised care. An algorithm can perform exceptionally well in a peer-reviewed validation study and still add little value in practice because the workflow around it assumes referral capacity, connectivity, or digital records that the deploying facility simply does not have.

Landry Dongmo Tsague, director of Africa CDC's Center for Primary Health Care, has repeatedly argued that AI cannot compensate for weak primary healthcare systems. Where facilities lack reliable electricity, connectivity, or functioning referral systems, AI does not solve the underlying problem.

A similar pattern is visible in diabetic retinopathy screening, a second area where African evidence has accumulated steadily rather than all at once. An early study using retinal images from Kenya's Nakuru eye disease cohort helped establish that automated detection was feasible in an African population. A more rigorous prospective validation followed in Zambia's Copperbelt province, where a deep-learning model (originally trained on a Singaporean cohort) was tested against retinal specialists' grading of images from over 1,500 Zambians with diabetes. It performed with clinically acceptable accuracy despite having been trained on a different population entirely.

Building on those results, a randomised controlled trial is now underway in Tanzania, testing not just diagnostic accuracy but whether AI screening actually gets more patients to timely ophthalmology follow-up compared with standard care. That progression, from feasibility, to cross-population

validation, to a trial measuring real referral outcomes, is closer to the ideal sequencing.

Yet even here, no African study has run the kind of independent, head-to-head comparison between competing retinopathy algorithms that exposed CAD4TB's relative standing. The gap that matters most—knowing not only whether a tool works, but how it compares with its alternatives on local populations—remains in even these better-evidenced cases.

Trust is earned through institutions

Popular commentary treats trust as something to be engineered into software or asserted through ethical principles. Ethical frameworks matter, but Africa is no longer starting from a blank page.

The African Union's Continental Artificial Intelligence Strategy identifies health as a priority sector while emphasising public value, local capacity, and responsible governance. Africa CDC has gone further through its Continental Health Data Governance Framework, building on the Malabo Convention and the AU Data Policy Framework to establish African approaches to collecting, sharing, and governing health data across borders.

The challenge today is less of drafting new principles and increasingly that of implementing, coordinating, and resourcing the institutions that already exist. Trust grows when ministries can independently evaluate products before procurement. It grows when regulators can monitor algorithms after deployment rather than only before approval. It grows when clinicians know that software updates, performance changes, and safety concerns are subject to continuous oversight. Trust is not an abstract ethical aspiration; it is an institutional capability.

Local representation is clinical evidence

The CAD4TB example illustrates a further lesson. Diagnostic performance is not a fixed property of an algorithm. It varies by population, imaging equipment, disease prevalence, and clinical setting—a phenomenon researchers call "dataset shift", in which algorithms lose performance when deployed on populations that differ from those on which they were developed and validated. That makes local representation more than a question of fairness; it makes it a question of clinical evidence.

Africa CDC has consistently argued that AI must be trained and continuously evaluated using data that is accurate, timely, and representative of the populations it serves, and that such data should remain governed within Africa, for Africa's benefit.

This is often discussed as data sovereignty. It is, as the TB and retinopathy examples show, a question of whether a health ministry can see how a given tool actually performs on its own population before relying on it, regardless of how sophisticated the underlying model appears on paper.

What success should look like

Africa should resist measuring progress by the number of AI tools deployed, or by impressive accuracy scores reported in validation studies conducted elsewhere.

We should be asking: can a TB programme in Kaduna or Kampala compare competing AI products on local data before procurement, rather than relying primarily on manufacturers' own studies? Can regulators continue evaluating those products as evidence evolves, just as the Tanzania retinopathy trial is now testing real referral outcomes? Will Africa CDC's Continental Health Data Governance Framework, once endorsed, be implemented and resourced at national level rather than joining the list of frameworks that simply remain on the shelf? Will the infrastructure ambitions set out in Africa CDC's Primary Health Care Digitalization Framework be realised so that AI tools are introduced into health systems capable of sustaining them? These are not questions about algorithms; they are questions about institutions.

None of this argues against the use of AI diagnostics in African health systems—quite the opposite. It argues that Africa's greatest contribution to the global AI conversation may not be developing the next breakthrough algorithm. It may be demonstrating how trustworthy AI is built: through resilient primary healthcare systems, strong public institutions, locally generated evidence, and continuous evaluation after deployment.

If Africa gets that sequencing right—infrastructure, governance, and locally validated evidence, in that order—trust will not be the obstacle to AI-enabled diagnostics. It will be its most durable outcome. ■

Re:

Caroline S Mbindyo is director of impact and innovation at Amref Health Africa. She is a futures researcher operating at the intersection of digital transformation, global health, and international development. With a strong foundation in innovation strategy, health R&D, and evidence-based practices, Caroline is dedicated to honouring local communities and is currently investigating how AI could shape health across Africa.

Data centre boom comes with steep health and climate costs

The need for minerals and materials as part of data centre infrastructure elucidates historical ties with imperial trade, conquest, and resource extraction in LMICs.
Photo: Arthouse Studio

Data centres are powering health advances at a rapid pace, but the industrial scale of their expansion threatens communities around the world. Behind AI's gleaming image sit growing concerns about the pollution and economic fallout of vast water and energy usage.

Ironically, it would take a high-functioning AI programme to find the vital connection uniting Kenya's Great Rift Valley, the culture-rich Canelones Department in Uruguay, and the rolling hills of northern Virginia in the US. All are steeped in natural beauty carved by millennia of geological development, but their current viral appeal is being fashioned by the scalpel edge of technology in the form of huge data centres powering the AI revolution.

The dependency-forming rapid response to our internet search and prompt requests—ranging from online consumerism to game-changing drug discovery—come at a price rarely considered by the everyday user. Each task undertaken by an AI model is routed through vast data centres packed with graphics processing units (GPUs) that generate heat, requiring energy and cooling on an industrial scale; one estimate put the electricity needed for the 5.1 trillion AI queries per year at 15 terawatt hours—the equivalent to the annual output of two typical nuclear reactors.

AI's image is that of a clean and gleaming technology but this is an industrial boom whose global data centre power usage is expected to increase to 219 gigawatts over the next five years—enough to power roughly 180 million homes. Observers believe low- and middle-income countries (LMICs) will bear the greatest environmental burden of this rapacious pursuit of computing capacity, with worrying flow-on consequences for health.

The people versus AI data centres

Campaigners are battling hard across Africa and South America to repel data centres that are packaged in an economic profile hard to resist even in energy-limited nations.

Microsoft's US\$1 billion project in Kenya's Rift Valley recently stalled because the government admitted the huge centre would suck up almost a third of the nation's energy generation before reluctantly suspending the project.

In drought-stricken Uruguay, a public protest moment was needed to persuade the government and Google to reshape its plans for a massive data centre near Montevideo after it was revealed it would siphon off vital drinking water supplies.

The issues for LMICs are compounded by the fact that many of these countries are the source of the raw materials used to manufacture processing chips, an industry that already causes environmental and healthcare damage. They are also often in dire need of the economic investment promised by big tech expansion plans.

Threat from data centres

The Earthkeepers vs AI Empires symposium in Lusaka, Zambia, in May, brought together global sustainability, environmental, and climate change experts and activists to frame the threat to LMICs from data centres. It stated: “The energy that powers these machines is drawn from landscapes where communities have been displaced by dams, mines, and vast development projects in service of feeding extractive economies. Even the cables carrying data follow paths shaped by centuries of imperial trade, conquest, and resource extraction.


“Across Africa in particular, territories once stripped for gold, diamonds, and oil are again being targeted; this time for the copper, cobalt, lithium, and other minerals needed to build and operate the AI empires.”

Enol Nieto, senior technician of governance for sustainability at the University of Central Catalunya, Barcelona, attended the landmark meeting and shared his experiences at the sharp end resisting Google’s Uruguay project that was cloaked in secrecy until he and Daniel Pena, a researcher and lecturer at the University of the Republic, Montevideo, challenged the government.

“Daniel spotted the data centre project on an Environment Ministry webpage and we started to investigate,” Nieto says. “I was living in Uruguay at the time, and it was a very difficult moment with the drought and losses of access to drinking water. There is human right to have access to water, so the data centre was in direct conflict with that.”

When government ministries failed to respond to queries, Nieto says he and Pena turned to public information requests and judicial proceedings with the support of a pro bono lawyer. The process revealed that Google’s planned data centre would consume an estimated 7.6 million litres of water a day. “This in a country with drought,” Nieto remarks.

According to Nieto, public disclosure of the figures sparked outcry and helped galvanise a social



movement around water. Opposition to the country's first proposed hyperscale data centre became one of the movement's central demands, with protesters taking to the streets for days. He says the public pressure pushed both the government and Google to revise the project, reducing its scale and replacing the planned water-based cooling system with electric cooling. Water will only be used for servicing bathrooms, kitchens, and fire safety on site.

“We should be able to discuss what model of data centres we want with transparency and all the information available. Europe is more advanced in its resistance processes and that may help the global south fight against exploitation.”

While Nieto describes the changes as a significant victory for water conservation, he argues the data centre will still place substantial demands on the country's energy supply—a challenge that is “everywhere”.



The global data centre supply chain relies heavily on industries in the global south such as those supplying raw materials for computer chips.
Photo: Heru Dharma

“We managed to change a project and to generate awareness about the materiality of the digital cloud on Uruguay society—that we see as a partial victory,” Nieto says, adding that in countries like Spain “they are targeting depopulated areas”.

“We are seeing developments in the Aragon region, in the north-east, that would use nine times the existing power used by the communities that live there. We are very concerned about the impacts on them and the environment.”

Nieto advocates for a global “moratorium” on the construction of data centres “until we know completely what the impacts will be”. “We should be able to discuss what model of data centres we want with transparency and all the information available. Europe is more advanced in its resistance processes and that may help the global south fight against exploitation,” he says.

A moratorium on data centres

The moratorium demand is echoed around the world. Friends of the Earth England, Wales and Northern Ireland campaigner Kierra Box highlights a jeopardy corroding communities, the environment, and climate change resistance: “The short-sighted scramble for AI is undermining our efforts to decarbonise and mount a global struggle against climate change.”

She says the “danger of data centres comes from their enormous energy and water usage”, explaining that a large data centre can use between four and 20 million litres of water per day—equal to consumption in 28,000–130,000 homes—and the same amount of electricity as half a million UK households annually.

“Some planned data centres may have good renewable energy proposals but most will either have on-site fossil fuel generators or require more oil and gas from the grid to meet their huge energy needs,” Box says. “There is far too little transparency at the moment [and] their expansion is outstripping regulation and political scrutiny. If we don’t get a handle on it, then [data centres] will proliferate and the environment will suffer.”

Box highlights that the environmental and social impacts of data centres extend far beyond the facilities themselves, especially in LMICs. “When we look at the longer supply chain feeding data centres, the raw materials that are in the computer chips are extracted from the global south—whose communities already bear the brunt of the impacts of pollution and human rights abuses we’ve known for years are linked to making batteries, microchips, and the infrastructure for data centres.”


Friends of the Earth groups around the world are raising concerns that governments, especially across Africa and Latin America, are struggling to balance environmental protections with the economic incentives offered by major technology investments. Box says this creates a risk that data centres could be developed in communities without adequate safeguards, amplifying existing environmental and public health issues.

“There should be a pause to think about why we are building data centres and who is benefitting. Only then should data centres be built more sustainably, with community consent, rather than chucking them in what are essentially in sacrifice zones,” she asserts.

Transparency needed to reduce health risks

The US—home to the world’s largest concentrated data centre hub in Virginia—is at the forefront of data centre expansion and the vanguard of health

In 2028, the public health impact of data centres could cause...



Source: Han, Y. et al., 2026, Health-Informed Computing: Estimating and Addressing the Public Health Impact of Data Centers.

concerns around their energy and water use along with light and noise pollution experienced by surrounding communities.

Research projects have highlighted healthcare concerns that noise and light pollution from data centres can cause hearing and sleep issues, while air pollution can increase rates of respiratory diseases and cardiovascular conditions.

An air pollution analysis by Harvard T. H. Chan School of Public Health concluded that emissions from the on-site gas turbines at one data centre in Virginia

could result in an estimated 3.4 to 6.5 additional premature deaths and increased hospital admissions and asthma-related outcomes.

A 2026 paper, which stressed that the public health impact of data centres is under-examined, found they they could cause 600,000 asthma symptom cases and 1,300 premature deaths along with other morbidity outcomes in 2028 if the data centre expansion rates continue at a high level.

Alison Steele, executive director of the Environmental Health Project, a non-profit public



Air pollution from data centres can increase rates of respiratory diseases and cardiovascular conditions, increasing hospital admissions and asthma-related outcomes.
Photo: cottonbro

health organisation in Pittsburgh, which has been identified as offering prime data centre development opportunities, believes that greater transparency from tech companies is a critical component of enabling the assessment and mitigation of health risks both in high-income nations and LMICs.

“We cannot just adopt wholesale a new technology, and one at this scale, without thinking through the public health consequences of it, or the equity issues that it creates,” she says.

“Data centres are continuously lit and produce a hum while they are operating, and research shows that low-frequency noise and light pollution contribute to stress levels, sleep disruptions, and mood disorders. There are added air quality issues from the fossil fuels powering data centres, either from their co-located power supply or the grid itself. There are health impacts along the entire supply chain and power generation journey that need to be considered.”

Local opposition

Local opposition has delayed or halted 75 data centre projects worth approximately \$130 billion in the US with opposition groups mobilising effectively across 49 states, according to a Data Centre Watch report. It states: “As the data centre pushback becomes mainstream, opposition is extending beyond grassroots groups in communities hosting new developments, drawing in statewide organisations, advocacy groups, and national voices focused on energy, water, land use, and ratepayer impacts.”

As the counter-data centre movement gathers pace in countries like the US, there will likely be concern that it will displace tech expansion plans to LMICs where companies might have an easier regulatory ride. But the global environmental campaigning community is determined to share intelligence, resources, and tactics to protect all under-threat regions and slow the data centre onslaught.

“This is a critical time for the global north to join forces with global south and mobilise around data centres.”

Major tech companies say they are listening and are committed to environmental sustainability. Google notes its “moonshot is to power our operations with carbon-free energy, every hour of every day, on every grid where we operate”. Microsoft says of its AI data centre buildout: “As we expand capacity, these systems are where many of our sustainability

commitments are put into practice—from reducing electricity, water, materials, and emissions impacts to working with partners on carbon-free electricity, grid resilience, value chain emissions, and carbon removal.”

Kierra Box, of Friends of the Earth, observes: “This is a critical time for the global north to join forces with global south and mobilise around data centres. We need to examine why we need these data centres, what they are being used for, and how to ensure they do good rather than harm.”

She advocates for greater transparency from tech companies and governments. “We need clear investigations into projected environmental and public health harms and these to be the first and foremost considerations. And we need decisions to be made with, not for, local people, shared so that we can build a system that benefits communities in all nations.”



Photo: Brett Sayles

How AI could crack vaccine science's hardest problem

The opinions expressed are those of the authors and do not necessarily reflect the position of Re:solve Global Health.

AI is set to help regional biotechs develop the next generation of vaccines by helping more researchers find answers faster and enabling new players to make the most of local knowledge.

A toddler in a village clinic in Malawi could, within a decade, get a rotavirus shot that works just as well as the one toddlers in London or Boston get every day. That is not the case today, despite rotavirus vaccines having existed for close to two decades. Anyone who works in this field knows the quiet embarrassment behind that gap: the vaccines we currently have tend to protect children in wealthy countries far better than they protect children in the low- and middle-income countries that carry most of the rotavirus disease burden.

Poor nutrition, maternal exposure, gut microbiome differences, competing infections—all of it chips away at how well a vaccine trains a still-developing immune system. Closing that protection gap costs hundreds of millions of dollars, and takes years we do not have to re-solve a problem we already thought we had solved.

Introducing the AI co-scientist

Much of that cost and delay stems from how vaccines are proven safe and effective in the first place. Normally, it is done the hard way: give people the shot, wait, and count how many still get sick compared to those who did not. That can take years. But for some diseases, researchers have found a more efficient way—a signal, often in the blood but sometimes in the stool or other bodily fluids, that reliably shows whether someone is protected, without waiting to find out the hard way.

Scientists call this a “correlate of protection”. It is why an updated Covid-19 or flu vaccine does not need a fresh multi-year trial every time a strain shifts. Regulators can clear it based on a much smaller study, one that only needs to show people had the

expected immune response. Correlates let vaccines keep pace with a mutating virus, bringing speed and cost savings without cutting corners on safety.

Many diseases do not have an established correlate of protection. Finding one usually takes years of systematic lab testing or a serendipitous discovery by a scientist in the field. But published scientific literature is now so vast that promising correlates may already sit buried in old papers, described once and abandoned for reasons unrelated to whether they worked.

Searching that volume of information would take a human team years. This is the promise of a new class of AI tools often called “AI co-scientist”. Used well, these tools can comb through decades of scattered research on rotavirus, malaria, or TB in hours, and return dozens of plausible candidates for new correlates of protection.

Benchmarking tools on the market

These new AI tools come with a risk. A genuinely new hypothesis has nothing to check itself against, unlike AI-generated clinical advice, which can be measured against existing guidelines. If we cannot tell how good a model's suggestions really are, a well-meaning researcher might pick up the wrong tool for the job, one that sends them down a dead end for months or years.

This is what PATH is working on, and we see our role as trying to prevent the bad outcomes we cannot afford. One part of that is benchmarking the tools already on the market, so a research team without a dedicated machine-learning group, or the resources to evaluate hundreds of candidate hypotheses themselves, knows in advance where a given tool holds up and where its limits lie. That can make the difference between a false lead caught early and one chased for months before anyone realises it was never real.

Getting a vaccine from an idea to an approved product can take 10 years and cost hundreds of

millions of dollars, most of it spent proving to regulators that the product works and is safe. A multinational company can spread that risk across dozens of programmes and shrug off the ones that fail. A regional biotech, closer to the patients but working with a fraction of the money, usually cannot survive even one failed trial. That is a large part of why so much vaccine development for diseases of poverty, such as TB and malaria, has relied on philanthropic and public funding rather than venture capital. The road was too long and too risky for anyone without deep reserves.

Towards a new era of access

If AI tools prove reliable, it will affect not only how quickly science moves, but also who gets to compete. Shorter, sharper trials mean less money is needed to get from candidate to approval, which changes who can participate. A company that could never have absorbed a decade of trial risk might only need to absorb one or two years of it, opening the field to teams that were priced out before.

None of this is guaranteed. Equitable access to these tools is its own problem, and solving it will take real investment. But if we get it right, the next rotavirus vaccine could be built by a team that understands local disease patterns rather than a global average. It could be brought to market by a company founded a few years ago in Nairobi, Hyderabad, or Dakar, close enough to the problem to see what a larger, more distant company would miss.

Go back to that toddler. AI alone will not have cured a disease; nobody serious would claim that. What will have changed is that a door once open to only the best-funded institutions is now open to many more. Getting through it will still take years of hard, unglamorous work. But there are potentially more hands able to try than there were even five years ago, and that is worth celebrating.

Vaccine-preventable diseases like rotavirus disproportionately affect children in low- and middle-income countries.
Photo: Adama Guebre



Re:

Dr Roberto Amato is a computational biologist with years of professional experience in AI and machine learning, and a career built in global health. He leads the AI for science portfolio at PATH, working to bring frontier AI tools safely and usefully into vaccine and drug development for diseases that have long been underserved.

Professor Bilal A Mateen is a physician by training and an internationally recognised expert in AI applications in health, spanning the implantation of frontier solutions in low-resource settings to biomedical science applications. He serves on the UN's Independent International Scientific Panel on AI and is the chief AI officer at PATH.

By Becky McCall

Better AI for women's health hinges on more data



Underfunding and a lack of available data in women's health research means that a diagnosis does not always lead to adequate treatment plans or professional support.
Photo: Polina Zimmerman

AI has the potential to advance women's healthcare, but only if it is built on higher quality data. Innovators are showing that closing longstanding research gaps and ensuring women's experiences are prioritised are the first steps towards more equitable care.

At first, receiving a diagnosis of endometriosis—a chronic and often debilitating gynaecological condition—felt like crossing an important threshold for Louise Dreisig. She expected the diagnosis to lead to the appropriate investigations, a coordinated treatment plan, and professional support. Instead, her battles were only beginning.

As a poorly understood and often overlooked condition affecting around 190 million women of reproductive age worldwide, endometriosis epitomises the longstanding inequality in women's health. Decades of underinvestment in research has left long-standing gaps in knowledge, diagnosis, and care, contributing to years of unnecessary pain and disability for many women. Common symptoms of endometriosis include severe pain during menstruation, chronic pelvic pain, and infertility.


Dreisig, now 48, from Copenhagen, Denmark, was diagnosed in her late 20s. She describes being bounced between hospital departments that failed to communicate, even as her condition deteriorated. Doctors appeared passive, she says, and she felt talked down to and dismissed. "I was often left to fend for myself and had to insist on getting help, all the while in severe pain," she recalls.

The absence of clear guidance left her feeling completely lost. Alongside the physical symptoms came what she describes as “mental chaos”, including questions about her future, her fertility, and whether the disease would affect every part of her life. “I was left alone with very significant existential questions,” she says, as well as “a deep sense of shame and exclusion from society”.

Dreisig's experience is all too common and helps explain why AI is generating such interest in women's health. AI has the potential to transform medicine by interrogating and integrating vast amounts of clinical, imaging, genomic, and patient-generated data, uncovering previously unrecognised patterns that could lead to earlier diagnosis, more personalised care, and even new biological insights. But its success depends on the availability and quality of those data—and this is the fundamental challenge to the use of AI in women's health.

Many conditions largely affecting women, including recurrent urinary tract infections and cardiovascular disease, as well as endometriosis, remain under-researched and hence lack the essential data needed. To further complicate matters, clinical information is often fragmented across specialties, diagnoses are delayed, symptoms poorly characterised, and women's lived experiences rarely captured alongside biological and clinical data. As a result, the rich, integrated datasets on which AI depends often do not exist.

The Lucy symptom-tracking app uses an algorithm to recognise health conditions like endometriosis and PMOS.
Photo: Condingo



AI-aided support across the diagnostic pathway

The EU-funded FEMaLe (Finding Endometriosis using Machine Learning) project is helping to address this deficit. Using AI to analyse vast amounts of patient information to identify subtle patterns that often elude traditional diagnostics, it set out to help doctors identify endometriosis earlier and with greater precision.

The project's findings helped raise the profile of women's health inequality in Denmark, contributing to the case for a new National Centre for Research in Women's Health. Its scientific and technological advances are also being carried forward through two major follow-on EU initiatives—EUMetrisosis and READI—while digital tools developed through the programme, including AI-assisted decision-support systems (SurgAR) and the Lucy symptom-tracking app, continue to evolve beyond the life of the project, which ended in 2025.

Progress depends not only on developing better algorithms but also on generating high-quality clinical, imaging, and patient-reported data.

Ulrik Bak Kirk, chief consultant and PhD faculty fellow in the Department of Public Health at Aarhus University, Denmark, who coordinated the project scientifically, believes FEMaLe highlights a broader lesson for women's health research. AI is not a shortcut to better diagnosis, he says, but one component of a much wider diagnostic pathway. Progress depends not only on developing better algorithms but also on generating high-quality clinical, imaging, and patient-reported data, understanding how diseases such as endometriosis vary between individuals, and integrating those insights into routine care.

He says FEMaLe pioneered AI-supported tools across the diagnostic pathway, including decision support for radiology and surgery, where machine learning helps clinicians identify and distinguish endometriosis lesions (when tissue similar to the lining of the uterus grows elsewhere in the body) while generating the structured data needed to refine future diagnosis and treatment.

Endometriosis is emblematic of wider inequalities in women's health, Kirk says, adding that symptoms are often normalised, unacknowledged, or attributed to other conditions. “I hate that it's stated as something benign,” he says. Although rarely life-threatening, “it reduces your quality of life. It affects your opportunity to get a proper education, to have a job...your everyday life.”

Conditions that primarily damage quality of life have historically attracted less research attention than diseases that threaten life. “You’re not going to die from it,” Kirk says of endometriosis, “but it will really cripple your life.”

Dreisig describes FEMaLe, for which she was an external advisor, as “an enormous gift” because it placed both the disease and those living with it firmly on the research agenda. She believes that the project exposed both the scale of the problem and how much work remains before women receive adequate physical and psychological care for endometriosis.

“If we are to move away from the shame of suffering from a gynaecological condition, it requires people who do not have endometriosis themselves to also take up the fight on our behalf,” she notes.

Amplifying unheard voices in women’s healthcare

For Saad Hussain, co-founder of Awaaz-e-Sehat (‘Voice of Health’ in Urdu), the research-led Pakistani maternal health platform, building AI for women’s health has been as much a lesson in humility as in technology. Designed to improve maternal healthcare in a country where more than 150 women die for every 100,000 live births, the AI-enabled platform, which launched in 2023, helps women generate and own their health records, access culturally relevant health information, and arrive at consultations better prepared.

Crucially, Awaaz-e-Sehat aims to place women, rather than technology, at the centre of care. After more than two decades in the technology startup sector, Hussain joined the project wanting to use AI for social good rather than purely commercial gain. What he did not anticipate was how much he would have to rethink his own assumptions.

“The most important thing was being flexible enough to understand that I still needed to learn. If I had built this product by myself, it would have been a disaster,” Hussain acknowledges, explaining that it succeeded because Dr Maryam Mustafa, an AI researcher at Lahore University of Management Sciences, led the work in close partnership with obstetrician and gynaecologist Professor Fozia Umber Qureshi, whose clinical expertise ensured the technology addressed patient needs.

Hussain’s reflection illustrates a wider lesson. AI cannot solve inequalities in women’s health if it is designed without women. Representation is not simply about including women in datasets, he argues, but about ensuring that women help define the problems technology is built to solve.



AI cannot solve inequalities in women’s health if it is designed without women from research stages to implementation.
Photo: Karola G

Overcoming the data deficit

Although working in very different healthcare settings, Awaaz-e-Sehat reached much the same conclusion as FEMaLe: before AI can improve women’s health, it must first overcome a more fundamental obstacle—the data simply is not there.

Realising that Pakistan’s public healthcare system lacked digitised maternal health data, the first challenge was to find a way to generate reliable, usable clinical information. The Awaaz-e-Sehat app was designed to capture women’s symptoms, medical history, pregnancy details, and investigation results in a structured digital format. “[The information] was sitting in paper files,” Hussain explains, adding that healthcare records remained fragmented across hospitals and provinces, with women often carrying paper notes from one appointment to the next. “Records were easily lost, tests repeated unnecessarily, and clinicians frequently lacked access to previous information.”

The first version of the app was primarily clinician-facing and aimed to facilitate consultations and identify serious complications, such as preeclampsia, at an earlier stage. The technology enabled spoken consultations in Urdu to be incorporated into

structured electronic medical records, while also creating a digital record, and through the algorithm, highlighting potential red flags and suggested follow-up questions based on established clinical guidelines. The system supported the creation of over 500 electronic medical records and flagged over 300 potential clinical risks during a seven-month deployment in a non-profit hospital.

“If we empower her to generate her own data and allow her to own that data digitally, the system has a much better chance of working.”

However, despite the benefits, overstretched clinicians struggled to use the app effectively. Accordingly, the team switched the emphasis from the clinician to the patient. “We realised that the only way this was going to work was for it to become more patient-focused. The one common element across the whole journey is the woman herself,” Hussain says, stressing that, “if we empower her to generate her own data and allow her to own that data digitally, the system has a much better chance of working.”

AI only works if it fits women’s lives

The latest version of the Awaaz-e-Sehat app enables women to record symptoms and other observations via WhatsApp in Urdu or English before their clinical appointment. They can also upload medical history and laboratory results, allowing clinicians to prepare before the consultation and make better use of limited appointment time. Designed around the realities of women’s lives and local healthcare delivery, these adaptations make the technology more practical and accessible, encouraging sustained use.

Women can also share records between providers using QR codes or web links, reducing repeat tests, retelling of their illness experience, and more efficient sharing of test results. “Laboratory results can sometimes be reviewed remotely, avoiding unnecessary journeys that, in many low-income households, require a husband or family member to miss a day’s work,” Hussain explains.

Crucially, the maternal health information provided via the app is also tailored to the Pakistani context rather than simply reproducing advice generated from Western healthcare settings. “It shouldn’t simply tell a woman to go for a run just because that’s common advice elsewhere in the world,” Hussain says. Likewise, advice is often tailored to be more relatable in the local setting. “Foetal growth may be described using familiar local fruits and vegetables,” he adds.

For Hussain, these decisions illustrate a wider principle for AI in women’s health. Algorithms trained predominantly on data from high-income countries cannot simply be transplanted into different healthcare systems and populations. “Whether AI reduces inequality depends on how you architect the solution and who you are building it for,” he says. “If you are building for a low- or middle-income environment but using data from a high-income country, then you are already on very loose footing. We are looking at the data available in Pakistan, generating the data we need, learning from it, and putting guardrails around the system to accommodate local needs.”

When AI is used to generate new evidence rather than simply analyse existing datasets, it has the potential not only to improve healthcare but also to narrow some of the inequalities that have long characterised women’s health.

His conclusion echoes a growing consensus emerging across women’s health research more widely. AI alone cannot compensate for decades of underinvestment, missing data, or neglected conditions. But when women themselves actively help shape its design, when technologies reflect local realities rather than imported assumptions, and when AI is used to generate new evidence rather than simply analyse existing datasets, it has the potential not only to improve healthcare but also to narrow some of the inequalities that have long characterised women’s health.

Dreisig, the Danish patient, says that is ultimately what AI should achieve: not simply faster algorithms, but recognition that women’s experiences are themselves a source of knowledge. FEMaLe, she says, “put the disease, and, not least, all those affected by it, on the agenda”.

More women are now coming forward to tell their stories, something Dreisig finds “deeply moving and incredibly powerful”. Yet, she says, this is only the beginning. “There are still so many doors that have yet to be opened, and even more questions that remain unanswered.” ■

Health AI must be designed with equity at its core

The opinions expressed are those of the author and do not necessarily reflect the position of Re:solve Global Health.

AI could transform African health systems, but without careful attention it may also amplify existing inequities. Equity must be built into how health AI is designed, validated, governed, and implemented.

AI is rapidly changing what is technically possible in health. Machine-learning models can detect patterns across complex datasets, support clinical and population-level decision making, improve diagnostic processes, and potentially extend scarce expertise to populations that health systems struggle to reach. Generative AI adds another dimension: the ability to synthesise and communicate health information at a scale we have not previously been able to manage.

For African health systems, these capabilities create significant opportunities. Yet my work across health research, implementation science, and data-driven decision making has made me equally interested in a more fundamental question: what happens when AI is introduced into health systems in which the underlying data, infrastructure, and distribution of power are already unequal?

This matters because AI does not eliminate the limitations of the systems from which it learns. It can reproduce them and, because algorithmic decisions can operate at enormous scale, potentially amplify them.

Inequity can begin with data

One of the most important lessons from my engagement with data science has been that algorithmic bias often begins before model development.

I encountered this through a health-financing initiative in East Africa that explored machine learning to estimate households' financial capacity. The policy objective was inherently equitable: to differentiate financial contributions while protecting households least able to pay.

But interrogation of the underlying data revealed a critical problem. Some of the most economically vulnerable populations were inadequately represented. A model trained on those data could therefore systematically overestimate the financial capacity of precisely the households the system was intended to protect.

This is an important distinction. A model may demonstrate good aggregate predictive performance and still distribute error inequitably across population subgroups. Overall accuracy therefore tells us relatively little about equity unless we also interrogate who is represented in the training data, where errors occur, and who bears the consequences of those errors.

The experience reinforced something increasingly recognised in AI governance: computational sophistication cannot compensate for information that is absent from the training data. The World Health Organization (WHO) has similarly cautioned that datasets that exclude rural communities and other disadvantaged groups can reproduce those exclusions in AI systems. For health systems, that is an equity problem, not merely a technical one.

Problem extends beyond algorithmic bias

Representation is only the first layer. Much of the discussion about responsible AI focuses on bias within algorithms. I am equally concerned about what surrounds them: data quality and provenance, the adequacy of validation datasets, weak governance structures, insufficient guardrails, and unclear accountability when algorithm-supported decisions cause harm.

Who decides whether a dataset is representative enough for a given use case? Are errors broken down by sex, geography, socioeconomic status and language, or only reported in aggregate? Somebody has to set acceptable thresholds for false positives and false negatives, and somebody has to watch for drift once a model is deployed. And when an algorithm shapes a consequential health or health-

financing decision, accountability has to sit somewhere concrete.

These are not problems that data scientists can solve alone. They require governance. WHO's framework for AI in health similarly emphasises human autonomy, transparency, explainability, responsibility, accountability, and inclusiveness, including mechanisms through which people adversely affected by algorithmic decisions can question those decisions and seek redress.

Health systems with interoperable digital records, reliable connectivity, computing infrastructure, strong regulatory institutions, and adequately resourced workforces will be better positioned to adopt AI than those without them.

These issues acquire another dimension in global health because AI is being introduced into a field already marked by significant asymmetries of technical, financial, and institutional power. African populations should not become primarily sources of data for technologies designed and owned elsewhere, while African health systems become downstream deployment environments with limited influence over how those technologies are designed, validated, and governed.

Data sovereignty is also about power

For this reason, we need a more ambitious understanding of data sovereignty. It cannot be reduced to the physical location of a server. It includes the authority to determine what data are collected, the purposes for which they can be used and reused, who has access to them, what health problems AI is prioritised to solve, and how the resulting scientific, social, and economic value is distributed.

This is why localisation cannot be an afterthought. AI systems developed predominantly using data, disease profiles, languages, and health-system assumptions from high-income settings cannot simply be presumed to generalise to African populations. External validation is important, but equitable AI requires more than validating imported models. It requires African researchers, health workers, communities, and policymakers to participate much earlier in defining the problem and shaping the technology.

Generative AI brings equity question closer to patients

These issues become particularly important as

generative AI moves from analytical applications into direct interaction with patients and communities.

Through my work in adolescent and women's health, I see enormous potential for AI-enabled digital tools to expand access to sexual, reproductive, maternal, and mental health information, particularly for young women who may face stigma, distance, cost, or confidentiality barriers when seeking care.

But generative AI changes the nature of risk. If a young woman seeks information about contraception, pregnancy complications, HIV, or mental health, the quality of the output depends on the evidence represented in the model, its ability to interpret her context and language, and the safeguards governing what it communicates. Incorrect, biased, or culturally inappropriate information does more than produce an inaccurate output; it can shape a decision with real consequences for her health.

WHO's guidance on generative AI for health therefore stresses the need for well-defined tasks, appropriate evidence, stakeholder involvement, and post-deployment auditing of large multimodal models. Human-centred AI means ensuring that technology augments rather than displaces human judgement, while protecting privacy, autonomy, empathy, and accountability.

Implementation science must accompany data science

There is one further dimension of equity: who gets to benefit from AI? Health systems with interoperable digital records, reliable connectivity, computing infrastructure, strong regulatory institutions, and adequately resourced workforces will be better positioned to adopt AI than those without them. Without deliberate investment, AI could therefore widen the very disparities we hope it will close.

For whom does [AI] work, and under what conditions? Can it be integrated into clinical and health-system workflows, and can staff act on what it recommends? And once it is deployed, who is watching its performance, and can it be governed and sustained at scale?

An algorithm can support an overstretched health worker, but it cannot repair a dysfunctional referral pathway. A diagnostic model is of limited use if the patient cannot get the test or treatment it

recommends. And no matter how sophisticated a prediction tool is, it cannot compensate for fragmented records or a work environment where staff cannot fit it into routine care.

This is where implementation science has to accompany data science. Asking whether a model works is not enough. For whom does it work, and under what conditions? Can it be integrated into clinical and health-system workflows, and can staff act on what it recommends? And once it is deployed, who is watching its performance, and can it be governed and sustained at scale?

Equity must be part of the architecture

AI presents an extraordinary opportunity for African health systems. We should not approach that opportunity defensively. But neither should equity become an ethical assessment conducted after the technology has already been designed.

Equity has to be built into every stage of the AI lifecycle: how problems are defined, how data are structured, how models are developed and validated, how they are implemented and monitored, and how accountability and redress are handled when something goes wrong.

African communities cannot simply be sources of data, and African health systems cannot simply be deployment sites. African scientists have a role earlier than that: not just validating a model built somewhere else, but helping define the problem it is meant to solve.

For me, every health AI system needs to be able to answer three questions: who is represented in it, who holds the decision-making power, and who is accountable when it fails.

Get that right from the start, and AI could help us build more equitable health systems. Get it wrong, and we will have used remarkably powerful technology to lock our existing inequities into the infrastructure of tomorrow's healthcare.

Re:


Dr Nyawira Gitahi, is a clinical researcher at the Kenya Medical Research Institute (KEMRI) and former Africa lead for DataDelta at IDinsight. She is faculty at KEMRI and holds an adjunct professorship at the University of Toronto. Nyawira currently heads the Africa-wide Leadership for Innovation and Excellence in Accelerating Research on Women's Health (LEA-WH) fellowship implemented with the National Academy of Medicine.

An algorithm can assist health workers in systems with strong digital infrastructure, regulation, and workforce capacity.
Photo: Soweto Graphics





Is AI closing care gaps for underserved communities?



AI has the potential for transformative applications beyond its use as an administrative tool in healthcare for marginalised communities.
Photo: Rendy Novantino

AI is reshaping how underserved communities access healthcare, from HIV testing to peer outreach. But early evidence of its effectiveness is mixed, and experts warn that bias and weak oversight could deepen the gaps it is meant to close.

Much has been made of the potential of AI to democratise healthcare, harnessing technology to fill in gaps in coverage, screening, and interventions for those who need it most. But, beyond the headlines, is this technology bringing needed healthcare to underserved and marginalised populations? Or is it mainly used as an administrative tool in these communities, while more transformative applications remain a goal for the future?

In countries like the US, there are concrete examples of AI delivering care and interventions to at-risk, marginalised communities. For conditions like HIV that disproportionately affect these populations, AI is already being used to identify who might benefit from testing or preventive treatment, and to connect people with care they might not otherwise seek out.

But scepticism of this tech persists, particularly around bias baked into the underlying data and a lack of clear governance around how it is used. Without careful oversight, AI risks reinforcing the disparities it is meant to solve.

Increase in use of AI tools

“AI tools are becoming increasingly mainstream in various aspects of the healthcare system,” says Dr Drishti Pillai, associate director of the Racial Equity and Health Policy Program and director of Immigrant Health

Policy at KFF, a leading US-based health policy organisation.

Pillai says that recent surveys reveal that a majority of hospitals and health insurers in the US report using AI tools, whether it is for filing health insurance claims, patient scheduling, medical coding, or administrative tasks. But outside of this health administrative space, among the public at large, there has been “a sharp increase in the use of AI tools for health”, she notes.

A disproportionate number of Americans turning to AI-powered tools to fulfil their healthcare needs are uninsured or from lower socioeconomic backgrounds.

In June, KFF released a nationally representative survey that revealed 31% of US adults reported turning to social media and 29% used AI chatbots at least once a month for healthcare information and advice, even though platforms like ChatGPT are prone to giving incorrect information. Within that number, one in five said they turned to these tech tools because they could not afford to access needed care. This was aligned with those who said “access and cost” were the drivers for turning to AI for health information in a separate KFF poll earlier in the year.

Dr Alex Billioux, chief health officer for Cityblock Health, a US-based provider delivering healthcare services to lower-income communities, says surveys like this underscore the reality that a disproportionate number of Americans turning to AI-powered tools to fulfil their healthcare needs are uninsured or from lower socioeconomic backgrounds.

Breaking that down even further, Billioux says Black and Hispanic people in the US are twice as likely as their white peers to turn to AI or large language model (LLM) tools to answer pressing healthcare questions.

He explains this is often due to historic mistrust of healthcare institutions among Black and brown communities in the US. He also attributes it to embedded systemic barriers to quality care that see people more willing to ask an AI chatbot a question than feel comfortable seeking a primary care physician—or even knowing how to access one in the first place. Lack of adequate health insurance coverage among low-income populations and communities of colour also makes AI tools attractive.

Billioux says that while it is encouraging to see these AI tools “democratise access to information”, the concerning reality is that only about 40–45% of people follow up on queries with healthcare professionals. This is especially discouraging in a



Mistrust of healthcare institutions and a lack of adequate health insurance coverage can drive people to use AI chatbots in place of primary care physicians.
Photo: Shvets Production

nation where recent figures show that the Black population experiences a 1.9 times higher mortality rate and the Latino population a 2.1 times higher mortality rate than their white peers.

AI in HIV interventions and care

One test case for AI's use in healthcare is HIV, which disproportionately affects people in marginalised communities. In the US, Black and Hispanic people made up 37% of new HIV infections in 2022, despite comprising 12% of the population. A 2025 global systemic review reports that AI tools have shown great potential in addressing existing gaps in HIV care, such as improving the accuracy of diagnoses and better treatment monitoring.

AI “has the potential to help scientists better understand how HIV persists in the body, identify potential drug targets, [and] predict drug resistance and treatment outcomes.”

“AI has been primarily used in prevention and care, where machine-learning tools are being used and studied to identify people who may benefit from HIV testing or pre-exposure prophylaxis (PrEP), predict treatment interruptions or loss to follow-up, and provide clinical decision support,” says Dr Andrea Gramatica, vice-president of research at amFAR, The Foundation for AIDS Research.

Gramatica says the clear value of AI is its ability to analyse “extremely large and complex datasets” and then identify patterns “that would be difficult for researchers to detect using traditional approaches”. He believes AI “has the potential to help scientists better understand how HIV persists in the body, identify potential drug targets, [and] predict drug resistance and treatment outcomes.”

He cites The HIV Immune Atlas Study, a three-year project that will unite leading HIV researchers and biomedical AI experts to create a comprehensive map of how HIV impacts the immune system and “how the viral reservoir differs across cells and tissues”. The goal is to harness AI to pinpoint the characteristics of cells that host HIV, hopefully creating models to predict where therapeutics may be used to “restore normal immune function and help eliminate the reservoir”.

From research to outreach, Gramatica says that evidence is still developing when it comes to how AI could be helpful. He says small studies show AI and automated interventions have helped raise awareness of, spread accurate information about, and increase uptake of PrEP, preventive medications that can lower risk of HIV infection by up to 99%.

Developers are also building “digital navigation tools” aimed at populations that have historically had less access to HIV prevention and care. “These provide private, multilingual information and connect individuals with testing, PrEP, treatment, or other services without requiring an initial visit to a traditional healthcare setting,” Gramatica says, noting that this kind of health data must be handled carefully. “Disclosure of HIV status, especially in certain geographies or cultural contexts, can expose individuals to stigma, discrimination, and, in some jurisdictions, legal consequences.”

AI finding the right messengers

Other researchers are using AI not to replace human contact, but to identify who within a community is best placed to deliver it. Dr Milind Tambe, the Gordon McKay Professor of Computer Science at Harvard University, worked with colleagues at the University of Southern California and Pennsylvania State University to create an AI system, called CHANGE (CompreHensive Adaptive Network samplinG for social influencE), that could pinpoint individuals in a social network who could promote accurate information about HIV prevention to their peers.

They tested this tool with more than 700 homeless young people who frequented three drop-in centres in Los Angeles, finding the method of combining AI-assisted social targeting and direct peer influence sharply reduced risk behaviours for HIV transmission.

“CHANGE showed that AI can be effectively deployed to reach and to assist improving health outcomes in vulnerable communities. This real-world study was highly encouraging,” Tambe says. He points to several follow-up projects inspired by these results, including one to improve maternal health in India. He and his colleagues took this work to start a collaboration with WHO for HIV prevention. He says the objective is to “use AI for effective recruitment for HIV testing.”

When asked if this unique algorithm to identify peer leaders among the homeless youth was more effective than traditional recruitment models, Tambe answers with a decided “yes”, explaining that this algorithm-centric approach was found to be more effective at ultimately reducing condomless anal sex —just one key HIV risk factor among this population—than more traditional, less tech-assisted recruitment and education methods.



Photo: Saad Ali

Addressing flaws and bias in AI

Despite the promise of AI in closing gaps in access to healthcare interventions for marginalised communities, Pillai notes a growing body of research that “suggests that sometimes a reliance on AI tools can exacerbate existing disparities in healthcare access”. Some studies, for example, have found that hospitals using AI tools to schedule patient appointments saw Black patients experiencing longer wait times than white patients.

Pillai says this may be due to bias—“the underlying data on which these tools are trained might be predominantly based on the data for white populations.” Or, she says, it might be because the AI tools correlate race with factors like socioeconomic status.

“There is potential for AI to be beneficial when used in healthcare, but the evidence so far is mixed. It really depends on how the AI is designed, what underlying models the AI is being trained on, and how it’s used.”

She adds that on the flip side there are findings that suggest, if designed correctly, AI could help “correct clinician bias”. In under-resourced hospitals, for example, having AI manage some of the administrative workload could reduce staff burnout and improve care for these populations.

“There is potential for AI to be beneficial when used in healthcare, but the evidence so far is mixed,” Pillai says. “It really depends on how the AI is designed, what underlying models the AI is being trained on, and how it’s used.”

One 2024 review concludes that to better address racial disparities in healthcare outcomes, there needs to be more regulations and ethical frameworks put in place in using AI tools. This includes bias detection tools and “active supervision by physicians”.

Pillai cites the Coalition for Health AI (CHAI), which has released guidance on the proper implementation of AI tools, and Encoding Equity Alliance, which is hoping to identify incorrect use of race in clinical algorithms and guidelines. “There is some movement to regulate how AI is used in healthcare with an eye towards equity,” she says.

Looking ahead to personalised medicine

Gramatica, of amFAR, The Foundation for AIDS Research, says that from the perspective of biomedical science, he expects and hopes that “AI will become part of the basic infrastructure of HIV research” in the next decade. He sees AI as an

“accelerator and a decision-support tool” rather than an autonomous scientist or clinician. He says there will be no replacement for the scientists and clinicians who will know how to wield and apply this data.

“With high-quality data, rigorous validation, transparency, privacy, and sustained participation from the communities using the technology, the long-term objective should be to give clinicians, scientists, and community-led organisations better tools to do their jobs,” Gramatica asserts.

Cityblock Health’s Billioux believes that in the coming era of more personalised medicine and intervention, AI has the potential to democratise to marginalised communities the same kind of “concierge” healthcare experience of which wealthy, privileged people in countries like the US already take advantage.

“Through the use of AI agents, medical professionals are able to synthesise all of this different activity that’s going on across their practice so that now you start to bring that level of personalisation and access to a population that otherwise wouldn’t be financially feasible to do,” he says.



AI is set to become an essential part of HIV research in the next decade, supporting rather than replacing scientists and clinicians.
Photo: Gustavo Fring

Why health AI needs trust built into workflows

The opinions expressed are those of the author and do not necessarily reflect the position of Re:solve Global Health.

In health systems in low- and middle-income countries, technology can travel faster than trust. Digital literacy, accountable follow-up, and local ownership decide whether AI becomes a trusted component of care.

Across our work designing AI-enabled primary care programmes with health system partners in South Asia, Africa, and in other low- and middle-income countries (LMICs), the same pattern keeps emerging. A conversation may begin with model accuracy, interoperability, or scale, but it soon moves to the questions that decide whether technology will be used: who explains the result? Can a frontline worker challenge it? Where does the data travel? Who owns follow-up when a patient leaves?

These questions matter where digitalisation lands inside systems already managing workforce shortages, fragmented records, variable connectivity, and long journeys to access care. But low resource must not be confused with low capability. Communities and frontline teams often know exactly where a workflow will break. Yet, too often, programmes seek their input only after a solution has already been deployed.

One field programme changed our question

Our largest completed field programme is Project Phoenix in Uttar Pradesh, India. It combines community screening, AI-assisted risk identification, telemedicine, and digital follow-up. Across more than 30 field camps, the programme screened over 10,000 people, trained more than 200 frontline workers, enabled over 1,800 calls with healthcare professionals, and reached more than 90% of its intended coverage.

Its research-led model also carried field learning beyond the programme. The Research Interest Score on ResearchGate—an indicator combining reads, citations, and recommendations—stood at 660.1; the work received more than 37,000 reads and reached audiences in over 74 countries. These

are not measures of clinical impact, but they show that the field lessons reached researchers and practitioners far beyond the programme. At the camps, the tool could flag within minutes that a participant's screening data might require clinical review. It was a prompt for a health worker, not a diagnosis.

The harder work began after the screen had done its job.

A screening flag created questions no model could answer. Would a family hear “risk” or “diagnosis”? Which clinician would receive the referral? Could the person travel? Who would call if they did not return? The people who made the technology usable were the frontline workers who translated a score into a conversation and knew which local clinician might respond. People trusted a person they already knew, who happened to be holding a digital tool.

Around 30% of participants entered follow-up pathways, including phone or in-person consultations with healthcare professionals, referrals for further assessment, and continued monitoring according to need. That figure shifted our question from how many risks a tool could identify to how reliably the system could carry a person from screening to care. The gap was not a footnote to implementation; it was the implementation.

Trust must be built within each health system

This lesson now shapes our proposed county-level pilot design in Kakamega, Kenya. We did not begin with a catalogue of AI products. We began by mapping how community health promoters, facilities, electronic medical records, referrals, and county decision-makers would connect across Lugari, Lurambi, and Butere. The objective is not three demonstrations; it is a locally owned care pathway that can survive three different service contexts.

Across our wider health-system engagements and programme development in Rwanda, Nigeria, Nepal, Sri Lanka, and Fiji, the starting point changes: community structures, languages, referral authority,



Digital health earns trust when it works where care happens, like this mobile health service for the PharmAccess mHealth programme in an African clinic. Photo: Elmvh, via Wikimedia Commons

and data rules differ. Scaling cannot mean copying an application; it means rebuilding a trustworthy workflow around the institutions people already use.

The Open Government Partnership 2023–2028 Strategy moves reform beyond consultation towards co-creation, implementation, monitoring, and accountability. Its digital governance challenge calls for public participation, digital inclusion, and oversight of AI. We are bringing that model into health: ministries, county teams, frontline workers, civil society, and communities shape the pathway, data rules, and measures of success before deployment.

Digital literacy belongs to the whole system

Digital literacy is often framed as a deficit within patients or frontline workers. Our experience suggests it has three layers: a patient needs to understand that a risk signal is not a final diagnosis, a health worker needs to understand the tool's limits and when to override or escalate, and the health system needs to know who receives the data, who acts next, and who remains accountable when the pathway breaks. A digitally confident community can still be lost inside a digitally incoherent system.

That changes training. Workers should practise explaining a result in the local language, managing uncertainty, recording consent, making a referral, and recovering after a failed connection. One-off training produces attendance; confidence requires supervised use, feedback, and repetition.

The current World Health Organization (WHO) Global Strategy on Digital Health 2020–2027 makes the systems point globally: digital initiatives need

financial, organisational, human, and technological resources. Its guidance on AI for health places accountability to health workers and affected communities at the centre of governance. That accountability must be visible when a result is handed over, not only in a policy document.

Make trust visible in the workflow

We therefore treat trust as a series of design decisions: language that separates screening from diagnosis; a named human who owns the next step; meaningful consent for screening, data sharing, and follow-up; a non-digital route for questions; escalation when a worker disagrees with the tool; and useful information returned to the frontline rather than extracted only for a distant dashboard. Overridden and missed referrals should become learning signals, not inconvenient noise.

Implementers, ministries, and funders should require a trust and digital literacy plan before deployment. It should specify who explains outputs in the local language, how people can ask questions or refuse data sharing, how workers challenge a result, and who owns follow-up. Its metrics should include comprehension, 30-day sustained frontline use, completed referrals, safe overrides, and equity across locations. These measures are harder to present than an accuracy score. That is precisely why they reveal whether health system digitalisation is working.

AI can cross a border in code, but care cannot. Care travels through local relationships, routines, and institutions. If we measure only what a model predicts, we will digitise the first mile and abandon the last. The goal is not to persuade LMIC communities to trust AI. It is to build health systems that behave in ways worthy of trust. ■

Re:

Aryan Chaudhary is director and chief scientific advisor at BioTech Sphere Research in India, a unit of NeoNexus Healthcare, and public governor (Rest of England) at Oxleas NHS Foundation Trust in the UK. His work spans responsible AI, digitally-enabled primary care, and implementation research across South Asia, Africa, the UK, and wider LMIC partnerships.

AI alone cannot fix Africa's health service delivery gaps



Health workers often choose tools that do not depend on internet connectivity over more advanced systems.
Photo: Gustavo Fring

AI offers hope for Africa's longstanding health service delivery challenges. But whether it makes good on its promise depends more on the strength of the systems built around it than the technology itself.

AI is being promoted as a way to stretch scarce health workers, predict supply shortages, and improve care. But emerging evidence from Africa suggests that its impact depends less on the sophistication of the algorithm than on the strength of the health system around it.

In some Tanzanian clinics, vaccination may take place in the open under a tree or in a room without a reliable power socket. That was the working environment Henry Mwanyika encountered while helping Tanzania's government and PATH develop an electronic immunisation registry. Rather than start with the most advanced technology available, the

team chose portable tablets and software that could continue working when internet connectivity failed.

"The solution itself has to be able to work without internet connectivity," says Mwanyika, then PATH's regional director for digital health in Africa. "So it has to be offline."

Indeed, AI is entering health systems where the biggest problems may lie beyond the technology: a missing nurse, an empty pharmacy shelf, an unreliable laboratory, or a referral that the patient cannot complete.

There are signs that AI can improve parts of care. Overcoming the wider service-delivery challenges shaping health outcomes across the continent is a far trickier proposition.

A real-world test in Kenya

One of the clearest and most recent tests so far took place across 16 Penda Health primary-care clinics in Kenya's Nairobi and Kiambu counties. Researchers embedded a generative-AI clinical decision-support

system into the clinics' electronic medical records. It analysed information entered during consultations and provided diagnostic and treatment guidance, while clinical officers remained free to accept, change, or ignore its recommendations.

The trial involved 103 clinical officers and 9,691 patients, and the results were mixed. Treatment failure—defined as a patient returning with unresolved symptoms, needing an unplanned escalation to higher-level or emergency care, or experiencing a safety problem such as a missed diagnosis or referral—within 14 days occurred in 2.2% of patients receiving AI-supported care and 2% of those receiving usual care, meaning the tool did not significantly improve service delivery in the two counties.

Better clinical reasoning may improve a consultation, but a patient may still be unable to afford medicine...or reach a hospital when referred.

But it did improve some parts of the care process. Clinical officers using it were more likely to document an appropriate diagnosis, produce a comprehensive note, and record an appropriate treatment plan. The system cost about US\$0.04 per consultation.

The findings show the limits of what decision support can change on its own. Better clinical reasoning may

improve a consultation, but a patient may still be unable to afford medicine, obtain a test, return for follow-up, or reach a hospital when referred.

The setting also matters. Penda Health is a private urban network with an established cloud-based record system and quality-improvement processes. The researchers cautioned that the findings may not transfer directly to rural clinics or less digitised public facilities, where patient needs may differ and issues like specialist clinician shortages and equipment shortages persist.

Dr Ikpeme Neto, founder and chief executive of the Nigerian health-technology company Wellahealth, says these surrounding constraints determine what clinicians can realistically do. "[In] Nigeria, even just doing a simple CT scan, you've got to be calculating and wondering and hoping," he says.

His company combines digital services with community pharmacies, point-of-care testing, and telemedicine support. Neto argues that hybrid systems are more realistic than digital-only care because patients still need trusted providers, tests, medicines, and referrals.

Data before algorithms

The setting up of Tanzania's electronic immunisation registry is a case study in why these foundations matter. The system linked individual vaccination records with stock management, allowing a dose recorded for a child to be deducted automatically from the facility's inventory.



AI might support a consultation, but a patient may still be unable to afford medicine, obtain a test, return for follow-up, or reach a hospital when needed.
Photo: RDNE Stock Project

A study involving 924 facilities found that average monthly vaccine stockouts fell from 7.1% before the system was introduced to 2.1% afterwards. It could not prove that the digital system alone caused the decline, but it suggested that more timely information could help facilities manage supplies. Other research found that health workers spent less time registering and vaccinating each child.

“[Health workers] have to generate data as a result of the work they do. They should not just be data collectors.”

But digitisation was not an uncomplicated success. Registry use declined over time in many facilities. Uptake was stronger where more workers had received training and clinics had stopped maintaining electronic and paper systems simultaneously. Connectivity, electricity, staffing, and supervision all affected continued use.

Mwanyika says the lesson was not simply to give health workers a device. The technology had to generate data through work they were already doing rather than turn clinicians into additional data clerks. “They have to generate data as a result of the work they do,” he explains. “They should not just be data collectors.”

Mwanyika describes change management as one of the programme’s most expensive components. Frontline workers, supervisors, government officials, and community representatives helped design and test the system, while Tanzanian companies were involved so it could be maintained after externally funded projects ended.

In this case, localisation meant local ownership of the workflow, technical capacity, and a government willing to continue investing.

Digital records for AI forecasts

Those digitally stored immunisation records later supplied data for a separate AI pilot. Technology company Macro-Eyes trained a machine-learning model using vaccine-use data from 710 Tanzanian health facilities. It forecast demand for individual facilities in two-week intervals.

In the two regions analysed, the model’s forecasts were substantially closer to actual vaccine use than the existing forecasting method. But accuracy was still uneven: at only around 45% of facilities did the estimate fall within the broad range of half to twice the number of vaccine doses actually used.

Macro-Eyes later said it planned to integrate and scale its AI-driven predictive system, but subsequent documentation by PATH described the technology being a pilot rather than a fully integrated national service.

The distinction matters for service delivery. An AI-based model may predict that a clinic will run out of measles vaccine, but it cannot release a budget, repair a vehicle, reopen an inaccessible road, or make doses appear in a regional warehouse.

Monitoring disease risk

Some service-delivery bottlenecks begin before a patient reaches a clinic and authorities need to know where disease risk is rising so scarce staff and supplies can be moved.

Dr Judy Omumbo, head of partnerships and resource mobilisation at the Science for Africa Foundation in Kenya, says African health systems hold potentially useful information, including historical data that were never formally recorded. That can include information from health facilities alongside population and environmental data.

The challenge is bringing those sources together quickly enough to identify emerging risks and act on them. “There are great data issues—how to integrate those data will be really important,” she says.

Timothy Endy, chief scientific officer at the Coalition for Epidemic Preparedness Innovations, describes a related use for AI-supported surveillance: combining outbreak information to identify emerging hotspots and help authorities decide where to send workers, vaccines, and other resources. He believes that such systems could “shift human resources to a hotspot that wasn’t anticipated”.

“I think [AI] can assist, but human beings are always going to be needed. It won’t be a solution on its own.”

But the value of the prediction still depends on whether trained staff, transport, supplies, and laboratories are available. Dr Wilmot James, senior adviser at Brown University’s Pandemic Center, similarly argues that earlier detection depends on the availability of clinics that can collect samples and laboratories capable of returning results quickly.

Assistance, not substitution

AI’s most useful contribution may therefore be relatively narrow: helping clinicians produce better notes, forecasting doses, identifying patients at risk of missing follow-up, or making specialist knowledge easier to retrieve.

Bermuda is much better resourced than most African health systems, but it shares a challenge common to small or underserved settings: limited local specialist capacity. Dr Chris Fosker, chief executive officer and medical director of Bermuda Cancer and Health Centre, uses algorithmic radiotherapy tools and remote peer review to draw on cancer specialists in Boston while treatment remains local. His preferred principle is to “move the knowledge” rather than the patient.

Fosker says safe adoption of AI requires capable information-technology teams, cybersecurity, clinical oversight, and a clear benefit for patients. Sometimes that benefit may simply be efficiency. “You’re gaining more time to be human,” he says.

But he warns against presenting AI as the answer to workforce shortages. “I think it can assist, but human beings are always going to be needed,” he says. “It won’t be a solution on its own.”

From pilots to public services

Technical performance is not the only obstacle between an AI pilot and a functioning public service. Endy says collaboration contracts, confidentiality agreements governing data and technical information, intellectual-property negotiations, and technology licences can take longer than parts of

scientific work. AI may accelerate analysis while deployment processes move at their previous speed.

Financing is another risk once political attention moves on. The World Health Organization’s (WHO) work on AI for health stresses that countries need governance, regulation, and implementation capacity alongside the technology. For health services, that also means funding training, connectivity, maintenance, data governance, and the people expected to act on AI recommendations.

For African health systems, the more useful question is not whether AI can overcome service-delivery challenges in the abstract. It is which bottleneck a tool addresses, whether local workers helped design it, whether it functions under real infrastructure conditions, and whether the health system can act on what it produces after the pilot ends.

An algorithm can help decide where attention is needed. The harder work remains ensuring that somebody, somewhere in the system, has the time, supplies, authority, and resources to respond. ■

For African health systems, the real test of AI is whether the technology works in real-world conditions and delivers results that can be sustained after a pilot.
Photo: Iwaria Inc



September 2026

About Re:solve

Re:solve Global Health is an independent journalistic platform for insights, conversations, and solutions to build more equitable health systems and healthier societies.

Despite continuous advancements in health, significant inequities and inefficiencies persist in global health systems.

With an international network of dedicated journalists, Re:solve Global Health delivers local and global insights that bring health challenges and solutions to the forefront.

We believe that real change begins when we empower communities, mobilise resources, and amplify the voices of those most affected by health inequities.

Re:solve Global Health is published by Re:Publications Group and owned by Dalberg Media Publishing. We are on a mission to rethink what is holding us back from living more sustainable lives, starting with how we resolve global health.

