PATH

PhAIx — Public Health Impact through AI

With the contribution of: OECD

AI and Public Health

Expert Dialogues on Strategy, Governance, and Ecosystem Transformation

Compendium of Virtual Expert Dialogue Series

June 2025 – November 2025

This compendium has been developed by PATH in collaboration with the OECD based on discussions and insights shared during the AI and Public Health: Expert Dialogues on Strategy, Governance, and Ecosystem Transformation series convened under the PhixAi initiative.

PATH would like to acknowledge the contributions of (in order of institutional hierarchy) Eric Sutherland, Resham Sethi, Aishwarya Rohatgi, Sarah Parwez and Deveshi Roy.

Contents

Foreword

Since ChatGPT was released in November 2022, health systems around the world have been thinking about how to apply artificial intelligence (AI) in their health systems. While it may have felt that the opportunities from AI were a shock to health systems, the reality is that, in many ways, AI is a 60-year-old mature technology. ChatGPT—and the birth of Large Language Models—represented a moment where the accessibility of the tool reached a threshold where it was available, affordable, useable, and potentially useful to everyone.

Having said that, the ability to leverage this powerful technology was only available to a small number of organizations that had the financial capability, data assets, and technical wherewithal to make use of it. This represented risk of expanding inequity divides experienced during the first wave of digitalization of health systems into untraversable digital chasms.

To mitigate expanding divides, the Organization for Economic Co-operation and Development (OECD) partnered with PATH through their Public Health Impact through AI (PhixAi) initiative. The goal of the partnership was to bridge the gap between the Global North and Global South to share experiences that would lead to more equitable and responsible use of AI across health systems.

This partnership led to a webinar series covering a wide range of topics including delivering AI in low-resource environments, considerations for the future of the workforce, and critical success factors for all countries starting on their AI journeys.

The webinar series was successful in reaching over 1,000 people over the six-month series. It brought together experts from Canada to South Africa, Denmark to Thailand, and Australia to India. This report attempts to summarize the hours of discussion and countless curious questions posed by the audience. I hope that the webinar series and this report demonstrates the effectiveness of multi-lateral partnerships that will be necessary in the years to come as the world guides the responsible scale of AI in health for the betterment of all.

It has been my honor to engage in this series with partners at PATH. I would like to express my gratitude for the leadership and friendship from Resham Sethi, Aishwarya Rohatgi, Deveshi Roy, and Sarah Parwez. I would also like to thank the over 30 individuals that helped bring this series to life. Their commitment and willingness to participate was humbling and inspiring.

I hope that the insights presented in this report will be helpful for everyone to take the necessary steps for the responsible development and scale of AI that generates better health outcomes for all.

Eric Sutherland
Organisation for Economic Co-operation and Development

Acknowledgment

The **AI and Public Health: Expert Dialogues on Strategy, Governance, and Ecosystem Transformation** series was established to convene leading thinkers spanning technology, policy, financing, and public health to explore the multifaceted intersections of artificial intelligence and global health. Delivered in collaboration with PATH and the Organization for Economic Co-operation and Development (OECD) under the Public Health Impact through AI (PhixAi) initiative, the series has provided a structured forum for rigorous dialogue, critical analysis, and cross-sectoral exchange.

This compendium consolidates the principal insights and thematic priorities that emerged throughout the series. It is intended to serve as a reference document for stakeholders seeking a comprehensive and nuanced understanding of the AI health ecosystem, including its institutional dynamics, systemic constraints, and the strategic considerations necessary to guide its responsible advancement.

I would like to extend my sincere appreciation to the distinguished experts and practitioners who generously contributed their time and knowledge to these dialogues. Your perspectives, grounded in real-world experience, have been invaluable in shaping the discourse, advancing collective understanding, and constructive dialogue on the role of AI in public health.

I am particularly grateful to the OECD team for their strategic guidance and steadfast support in the development of both the dialogue series and this compendium. Their commitment has been integral to elevating this initiative and amplifying its impact.

Finally, I express my deep gratitude to the team at PATH for their diligence, and sustained efforts in executing this initiative.

Sameer Kanwar
Director, Digital Health, AI and MedTech, PATH

List of Abbreviations

AI	Artificial Intelligence
API	Application Programming Interface
CDSS	Clinical Decision Support System
CHW	Community Health Worker
CSR	Corporate Social Responsibility
EHRs	Electronic Health Records
LMICs	Low- and Middle-Income Countries
ML	Machine Learning
OECD	Organization for Economic Co-operation and Development
PhixAi	Public Health Impact through AI
RCTs	Randomized Controlled Trials
TGA	Therapeutic Goods Administration
WHO	World Health Organization

Introduction

Artificial intelligence (AI) has moved from experimentation to implementation across global health systems. Over the past five years, AI tools have been deployed for scribing health encounters, diagnostics, surveillance, and decision support. Yet the pattern across regions remains the same—rapid innovation at the edge, slow institutional integration at the core. Most initiatives remain fragmented, pilot-based, and donor-driven, with limited evidence of long-term adoption or national ownership.

The **Public Health Impact through AI** (**PhixAi**) initiative is PATH's flagship program to shape the global AI-for-health ecosystem by translating AI from potential to measurable public health impact. Anchored in the pursuit of universal health coverage and equity, PhixAi connects innovators, policymakers, and implementers to build AI systems that are safe, inclusive, and sustainable.

As part of this initiative, **PATH** and the **Organization for Economic Co-operation and Development** (**OECD**) co-hosted the **AI and Public Health: Expert Dialogues on Strategy, Governance, and Ecosystem Transformation series** between June and November 2025. The six-month virtual series was designed to advance collaboration and leadership on the responsible and equitable use of AI in health systems. Each 90-minute session examined a key structural dimension, from designing AI for low-connectivity environments and scaling use cases within public systems, to financing sustainable adoption, strengthening the health workforce, and advancing regulation, ethics, and global governance.

Developed from these discussions, this report synthesizes insights from global experts and institutions working to advance system-ready AI adoption in public health. It serves as a concise knowledge product that outlines practical pathways for integrating AI sustainably and equitably into national health systems. Table 1 provides an overview of the sessions.

Table 1: **AI and Public Health: Expert Dialogues on Strategy, Governance, and Ecosystem Transformation series (June–November 2025) at a glance.**

Session	Title	Core Focus
1	Designing AI for Low-Connectivity and Resource-constrained Health Systems (**June 2025**)	Adapting AI tools to function effectively in low-resource and offline environments
2	From Use Case to Scale Culture: Embedding AI in Public Sector Practice (**July 2025**)	Moving from isolated pilots to system-wide integration
3	Financing the AI Ecosystem for Health (**August 2025**)	Building sustainable financing models for system integration
4	AI and the Invisible Workforce: Task Shifting and Automation Risks (**September 2025**)	Understanding AI's impact on the health workforce and digital labor (with a focus on youth)
5	Regulation and Ethics: From Guidelines to Enforcement (**October 2025**)	Translating AI ethics into enforceable, context-specific governance frameworks
6	Global Health Diplomacy and the AI Frontier (**November 2025**)	Examining AI's role in reshaping multilateral health governance

Together, these discussions formed a comprehensive inquiry into how AI can be designed, financed, governed, and institutionalized in ways that strengthen equity and long-term system performance. Ultimately, the series pointed out a way to balance market, health, and sustainability forces to enable practices where AI in Health can scale fast while doing no harm.

Objectives of the dialogue series

The series was designed to:

- Identify the operational and institutional enablers to integrate and responsibly scale AI solutions.
- Examine financing pathways that extend beyond tool procurement to include data systems, infrastructure, and workforce capacity.
- Strengthen regulatory and governance frameworks to ensure accountability, equity, and transparency.
- Advance cross-sector understanding between public health, primary care, clinical research, innovation, regulatory, and policy actors.

Context and rationale

AI's entry into health systems coincides with unprecedented strain on global public health infrastructure. The World Health Organization (WHO) projects a shortage of 11 million health workers by 2030, predominantly in low- and middle-income countries (LMICs). At the same time, more than half of health facilities in LMICs face intermittent power and connectivity, limiting the feasibility of cloud-dependent AI systems. Most digital health platforms remain fragmented, with parallel data collection portals, incompatible standards, and minimal interoperability.

AI applications in health have shown encouraging outcomes across diverse contexts, including tuberculosis (TB) screening, disease surveillance, and clinical documaentation. Studies evidence that over 80 percent of AI solutions show benefits in randomized clinical trials (RCTs). At the same time, many of these initiatives remain at an early or pilot stage, with challenges related to integration, financing, and governance still under discussion. Simply put, innovators are making progress in developing innovative solutions; however, the governance, policy, process, data, and technical environment is not suitable to efficiently scale AI solutions to enable benefits to reach every person in every location.

Funding often centers on initial innovation rather than long-term adoption, and procurement systems are only beginning to adapt to software-based and algorithmic

> As countries move toward more coordinated approaches, ensuring equitable access, institutional readiness, and clear accountability will remain central to achieving meaningful and sustainable AI integration in health systems.

tools that support co-operation across health systems. Regulatory pathways for AI in health are evolving, but gaps remain in evaluation, validation, approval, and post-market monitoring processes. These factors contribute to continued reliance on external support and limit the transition of pilots into sustained national programs. Efforts to strengthen evidence, build local capacity, and align financing with system needs are ongoing. Ethical and governance frameworks are also advancing globally, though implementation and enforcement vary across contexts.

As countries move toward more coordinated approaches, ensuring equitable access, institutional readiness, and clear accountability will remain central to achieving meaningful and sustainable AI integration in health systems. To support economic co-operation, these approaches will benefit from some cross-border collaboration to simplify the ability to responsibly scale AI solutions. Such an approach would prioritize long-term sustainability and benefit while supporting short-term objectives.

Purpose of this report

This consolidated report synthesizes key insights from all six dialogues and serves as a resource for decision-makers across government, donor, and implementation ecosystems seeking to integrate AI responsibly into health systems worldwide. The report focuses on system-level design rather than individual technologies, highlighting recurring themes such as co-design, localization, predictable financing, ethical accountability, and workforce preparedness.

Subsequent sections titled webinar 1 to 6 align directly with the six webinars convened under the **AI and Public Health: Expert Dialogues on Strategy, Governance, and Ecosystem Transformation series**. Each section distils the dialogue’s central discussion points and captures its core challenges, implementation mechanisms, and the stakeholder perspectives that emerged. The penultimate section consolidates the key recommendations gathered across the dialogue series into a structured, stakeholder-wise framework, tailoring guidance to the roles and responsibilities of each relevant stakeholder.

Webinar 1

Designing AI for low-connectivity and resource-constrained health systems

Co-authors: Anne Muthoni Gitimu | Sameer Pujari | Arnab Pal | Gwydion Williams | Dr. Shantanu Pathak | Eric Sutherland Resham Sethi | Deveshi Roy | Morgan Lemin

While many initial breakthroughs for the use of AI occurred in well-financed and technically capable environments, most of the world lives in areas of low-connectivity or resource-constrained health systems. While efforts are being made to improve connectivity, the reality remains that the resources needed to develop, deploy, and use AI are significant and are only available for a limited number of stakeholders.

In Webinar 1, panelists shared several examples and experiences that highlighted the value of collaboration, leverage, and ingenuity to address systemic gaps and barriers.

Key discussion points from Webinar 1 included:

User-centered innovation: Building trust and designing with context in mind, particularly in challenging environments.

→ **Resource-constrained implementation:** Operating efficiently despite constraints related to data, financing and capacity.

→ **Intentional transition to scale:** Acting with intentionality to transition from design to scale and integration.

This chapter will provide some of the practical methods that are being used to make effective use of AI despite limitations.

1. Centering innovation on users in challenging environments

Trust, stigma, and frontline realities

Introducing AI in public health highlights significant user-level concerns and behavioral barriers. One major challenge lies in fear and stigma. For instance, when AI-powered chest X-ray tools for TB screening were introduced in slums and rural areas, health care providers were initially hesitant to trust the software's diagnostic ability. Moreover, community members expressed concerns about stigma due to reports showing marked lung abnormalities before a formal diagnosis.

These experiences underscore the necessity of human-centered design that minimizes unintended harm, including stigma. Collaborative efforts with technology providers were critical to adapting tools to local sensitivities. Ultimately, adoption hinges on building user trust by ensuring that tools are reliable, transparent, and designed to respect community perceptions.

Designing for use: testing, workflows, and iteration

Designing AI tools for primary care requires strong focus on explainability, accountability, and structured workflows. Tool adoption varies based on users' technical proficiency, perceived value, and task alignment.

Tools should simplify rather than add burden. Working with community health workers (CHWs) revealed that overloading them with additional tasks undermines tool adoption. Tools must aid, not complicate, their workflows. For example, some GenAI models were found to flag all test cases as positive, maximizing true positives but significantly increasing false positives, thereby undermining diagnostic utility. In cases where tools flagged every case as high-risk, CHWs became overwhelmed without a clear pathway for action.

> "One lesson from foetal monitoring tools was the importance of making the outputs traceable. Health workers need to understand and explain why a patient was flagged high-risk."

Moreover, health workers may reject tools that feel like an additional burden or fall outside their role. Poorly integrated solutions (e.g., TB cough detection requiring repeated trials) can frustrate users. Multi-device setups—such as separate laptops for image capture and AI interpretation—also increased complexity and breakage risk. **Seamless, single-platform integration emerged as a key principle for ease of use.** Solutions that run in the background (e.g., AI-powered X-rays) were well received due to their seamless integration.

BOX 1 CHATBOT INTEGRATION IN KENYA

Integrating chatbots targeting youth in Kenya succeeded because of high mobile phone penetration and youth familiarity with digital tools. In contrast, solutions like ultrasound screening, which require integration with public infrastructure, faced greater challenges.

Workflow integrity must be preserved, and efforts must align with CHWs' existing duties, enabling them to triage without requiring them to make decisions outside their clinical authority. CHWs can triage, but decision making should rest with medical officers. This design distinction is vital in systems with rigid hierarchies. Pilots demonstrated that unless AI tools align with the decision-making protocols of public health systems, they face limited adoption.

BOX 2 CLINICAL DECISION SUPPORT SYSTEMS: LESSONS LEARNT FROM DEPLOYMENT

The deployment of Clinical Decision Support Systems (CDSS) in Kenya, Nigeria, and Rwanda illustrates the importance of user testing in early phases. In one pilot project within primary health care systems in Kenya, CDSS was initially designed to support doctors by integrating a chatbot feature into the clinical interface. This chatbot was accessible through a small icon that allowed doctors to ask questions and receive guidance while logging patient notes. However, the feature was not utilized, as clinicians, already managing heavy patient loads were reluctant to engage with a system they were unfamiliar with or did not yet trust.

This experience led to a key insight: AI tools in clinical settings cannot rely on proactive engagement by overburdened health care workers. As a result, the design was refined to create an "always-on" system that automatically generated feedback after a doctor completed patient notes. The updated interface provided a traffic-light–based risk assessment (red, amber, green), along with recommended care actions, requiring no additional user input.

Further lessons included alarm fatigue from over-flagging high-risk cases and varying needs for localization. For experienced clinicians, localization was less critical, as they viewed the tool as a secondary guide. In contrast, for CHWs, localized and actionable outputs were essential. These insights stress the need to test assumptions through iterative design and prioritize field-based testing.

2. Operating efficiently with constrained resources

Data limitations

Local data limitations, whether in access, quality, or structure, significantly impact the scalability of AI solutions in public health systems. In many cases, public datasets are vast in volume but lack depth in the specific features required for solving a targeted health problem. For example, a dataset may include millions of entries, yet only a fraction may be usable when training models for a precise task. This limits the ability to develop robust models or move beyond the proof-of-concept phase.

A central challenge lies in the lack of generalizability. Even when AI tools are successfully piloted in a specific population or region, replicating their performance across diverse geographies, such as across India's own sub-national variability, poses risks. To address this, there is a pressing need for a framework that enables validation and performance assessment across varied populations. Such a framework would help ensure both the localization of tools and the confidence of decision-makers in their reliability.

> “Replicating AI across different populations is a challenge. We need validation frameworks that give decision-makers confidence in their own context.”

There is a distinction between solutions developed for controlled environments—such as hardware-integrated tools for fetal monitoring or pathology—and those intended for system-wide use. Closed-loop systems with integrated learning and software-hardware feedback loops are more amenable to scale. However, AI tools designed for general-purpose use or as digital public goods face higher risks due to the variability of real-world data environments. Tools validated in controlled settings often perform inconsistently when deployed in different demographic or epidemiological contexts.

BOX 3 GENERATIVE AI VERSUS TRADITIONAL MACHINE LEARNING IN LOW-RESOURCE HEALTH SETTINGS

When assessing the suitability of traditional machine learning (ML) versus GenAI in low-resource health environments, three key dimensions stand out: cost, performance, and practicality.

“Generative AI is orders of magnitude more expensive than traditional ML.
The models are so large they must be hosted remotely, driving up cloud and compute costs—more complex and more expensive overall.”

Recent comparisons between multimodal GenAI models and traditional computer vision models—particularly in the context of rapid diagnostic testing—reveal computing performance challenges. GenAI systems often suffer from reliability issues such as application programming interface (API) throttling, where outages or rate limits delay or prevent real-time access to model outputs. Furthermore, if the model provider changes the system behind the API without notice, it creates downstream risks for any health system depending on its output. Traditional ML, in contrast, offers better transparency and stability, making it more practical for regulatory approval and consistent use in public health systems.

While GenAI offers remarkable flexibility—especially for clinical decision support, where a broad domain understanding is advantageous—it remains less suited for narrow, high-accuracy diagnostic tasks unless fine-tuned and optimized at smaller model scales. Until infrastructure, regulation, and cost barriers are resolved, traditional ML remains the more viable choice in most low-resource health settings.

Financing, capacity, and ecosystem

Governments often face high switching costs in adopting AI tools, especially when legacy systems are already in place. Consequently, **while pilots may show promise, scaled adoption demands more than technical proof**—it requires financial feasibility, robust data infrastructure, and long-term support for integration. These conditions are not yet met across most low-resource public health settings. GenAI is significantly more expensive than traditional ML, primarily due to the size of the models and the requirement for remote cloud hosting. This infrastructure adds complexity and cost, particularly in LMICs.

Even if AI improves screening accuracy, adoption is unlikely unless the cost of AI-enabled tools is comparable to existing systems. Business models that combine hardware and software components also raise complexity in procurement and budgeting. Communication strategies should demystify AI for nontechnical stakeholders, **articulating cost savings and health benefits clearly.**

BOX 4 THE HIDDEN COST OF LANGUAGE

An important, often overlooked cost driver is language. In non-English contexts, tokenization processes generate more tokens on average, increasing both computational and financial costs. For instance, in Rwanda, questions posed in Kinyarwanda were found to be more expensive to tokenize and process than the same queries in English. This underscores the need to design cost-conscious, language-inclusive models for LMIC deployment.

Successful integration depends on both financing and capacity. Solutions piloted in collaboration with the private sector—such as predictive nutrition models using weather and health data—show promise. Incentive structures, including financial or non-monetary rewards, can enhance CHW engagement. Importantly, financing strategies must be defined early on. Whether through government budgets, donor funds, or user fees, sustainable financing is necessary for scaling. Clarity is needed regarding financing mechanisms (e.g., social insurance, out-of-pocket payments, and ongoing maintenance costs).

3. Acting with intentionality to transition from design to scale and integration

Transitioning AI tools from pilot stages to large-scale health system integration requires both technical and institutional foundations. While the technical architecture—such as model development and algorithm design—is evolving rapidly and remains relatively system-agnostic, institutional factors are more foundational and context-dependent. Most digital health systems in LMICs are fragmented and not ready for AI at scale. Policymakers face a dilemma: invest in full system interoperability or support integration-friendly tools in the short term. Solutions like Newton's Tree and OpenFn offer "integration-as-a-service" models to help overcome current ecosystem fragmentation. In the short term, supporting such tools can enable deployment without requiring full ecosystem reform.

Three critical elements underpin the transition to scale:

1. **Foundational enablers:** Many LMICs lack robust governance frameworks for AI regulation, ethics, and evaluation. Effective scale-up begins with robust country-level governance structures, clear benchmarking standards (e.g., software specifications and content standards), and localized (linguistic, cultural, and clinical) deployment strategies. These are essential for enabling interoperability, ensuring ethical oversight, and validating AI performance in varied contexts. GenAI presents notable challenges around regulatory approval. The underlying model architecture and parameters are typically hidden (i.e., software of unknown provenance), regulatory agencies may struggle to validate or certify the tool as a medical device.

2

Problem-solution alignment: A key priority is identifying high-impact, systemic pain points—such as gaps in the health workforce or the administrative burden on primary care providers—and building AI tools that alleviate rather than add to these challenges. Many AI solutions fail to scale due to lack of alignment with national infrastructure, scale-up planning, and early engagement with public systems. Past digital health efforts also often failed because new tools imposed additional learning curves or tasks on frontline workers. Newer AI systems must be designed to integrate into existing workflows and simplify processes.

3

Sustainable investment and market shaping: A persistent challenge is the gap between research, innovation, and scale. Many promising tools remain in pilot phases due to limited engagement from donors, governments, and procurement systems. Sustainable investment strategies must involve not just the technology creators and implementers, but also the broader ecosystem—including funders, regulators, and ministries of health. Market shaping requires aligning donor priorities with local system needs and ensuring there is a pathway from innovation to procurement.

> “Do we wait for a transformation to fully interoperable systems, or do we act now with what we have? Integration as a service is our near-term lifeline. Without it, innovators keep solving the same problems over and over again.”

In summary: the power of innovation

Scan here to watch the full webinar:

A major insight emerging from frontline implementation is the transformative impact of AI-powered tools, especially when used in conjunction with point-of-care ultrasound (POCUS). In many African communities, these tools have enabled access to timely and quality antenatal care, especially for high-risk pregnancies. AI-driven diagnostics play a critical role in reducing maternal mortality and ensuring consistent follow-up care.

While early AI breakthroughs largely emerged from well financed and technology abundant environments, this need not define the future of AI in health. Purposeful, context driven approaches to design, deployment, integration, and scale can enable AI tools to function effectively in low connectivity and resource constrained health systems. When innovation is undertaken with this intention, it can extend beyond a limited group of stakeholders and support equitable last mile delivery.

Webinar 2

From use case to scale: embedding AI in public health practice

Co-authors: David Hansen | Susanna Aussó | Scott Hayden Mahoney | Enya Séguin | Victor S. Lubuku | Eric Sutherland Resham Sethi | Deveshi Roy

Integrating AI into public health practice is important because it shifts AI from being just a technical innovation to becoming something that improves population health, equity, and health system decision-making in real-world settings. The rapid rise of AI in health systems has sometimes led to "pilotitis," where many small pilot projects show promising results but fail to scale. Although these initiatives demonstrate the potential value of AI, they often remain at the testing stage due to barriers such as limited funding, weak infrastructure, and difficulties integrating into existing health systems. As a result, successful innovations do not always become part of routine public health practice. This can also risk wasting scarce resources on duplicative investments. In addition, unmanaged pilots may fragment personal health data across different facilities or applications instead of maintaining complete, integrated records for individuals. As a result, successful innovations do not always become part of routine public health practice.

This webinar explored the issues of scale—of how solutions are initially evaluated to establishing environments that encourage moving from pilot to scale by design.

Key discussion points:

- **Having a systems approach to scaling AI** by embedding AI into health policy and service delivery, creating multi-level institutional pathways beyond pilots, and strengthening digital readiness.
- **Ensuring the validity and contextual fit of AI tools** through robust evaluation frameworks and deep localization.
- **Designing strategies that enable system-wide scale** and foster a culture of ownership and sustainability to ensure lasting impact.

1. Having a system's approach to scaling AI

Embedding AI into health policy and service delivery

A systems approach is essential for scaling AI in health care, revealing two parallel strategies: laying strong foundations and building adaptive institutional mechanisms. The goal is not to deploy isolated tools but to embed AI into governance, infrastructure, and the culture of health care delivery.

In Australia, the government has taken early steps through regulatory clarity and digital infrastructure. A national AI Transparency Statement signals intent, though it currently applies more to internal use than clinical practice. More developed is the framework for software as a medical device, which provides a legal backbone for AI adoption. State-level health systems are also experimenting with guidelines for real-world use, such as AI scribes in hospitals and primary care.

Nationally, the focus has been on data standards and interoperability to ensure electronic health records and international patient summaries are in place before widespread deployment. The lesson: embedding AI begins not with algorithms, but with data, regulation, and trust.

In Catalonia, a more programmatic approach has created mechanisms that move innovations from pilots into routine practice. Public "AI Challenges" are designed around concrete system needs, ensuring solutions address real gaps. Bottom-up projects from hospitals and research centers receive support to evolve into system-wide tools. An AI observatory tracks over 180 initiatives, fostering synergies and reducing duplication. The lessons are clear: start with needs, not technology; build scalability and sustainability from day one; adapt procurement and governance to the pace of innovation, and ensure clinical leadership so adoption is shaped by the realities of care delivery.

Together, these examples show a shared insight. Successful AI integration is less about any individual tool, and more about the systems, institutions, and readiness that surround it.

Creating institutional pathways beyond pilots

Scaling AI beyond pilots remains one of the hardest challenges in public health.
Breaking the cycle requires early and frequent engagement that must happen simultaneously at three levels: (i) government leadership for policy support and resource allocation, (ii) ministry program managers for integration planning and budget alignment, and (iii) frontline health workers for workflow adaptation and practical usability. **Engaging at only one level inevitably leads to failure: policy without frontline buy-in, or end-user enthusiasm without ministry support, cannot sustain scale.** When users are involved in design, they build ownership and resistance diminishes.

However, engagement alone is not enough. Sustained adoption depends on ecosystems: governments, academia, private innovators, civil society, and donors working in alignment. Without clear governance and sustainability models in place from the start, governments hesitate to commit, knowing donor funding may disappear.

Momentum is building in countries energized by AI's potential, such as Rwanda, Kenya, and Senegal. Dedicated scaling hubs are being established in Rwanda and Nigeria to identify high-impact tools and embed them within government-led systems. **Thus, for AI to move from pilot to practice, the conditions for continuity must be built from day one.**

Strengthening digital readiness for scalable investment
Before AI can deliver on its promise, the digital foundations of health systems must be in place. Reliable internet, consistent power supply, affordable hardware, and basic digital literacy remain fragile in many settings, and without them, even the most promising tools collapse in practice.

Beyond the basics, the maturity of health information systems matters. Effective AI requires data that is not only available, but interoperable, yet in many LMICs, health data remains fragmented across incompatible systems, driving up costs and limiting scalability. Infrastructure constraints compound this further: unreliable connectivity and frequent power outages are common realities in these settings. Solutions must therefore be designed for resilience from the outset, functioning even when systems go offline.

Readiness is not only about systems but about people. AI adoption is a process of behavior change. With a projected global shortage of 11 million health workers by 2030, the question is not just how to support the existing workforce, but how to prepare the next generation for an AI-enabled world. Integrating digital and AI literacy into health education is as important as deploying the tools themselves.

> “AI scale-up is really deeply tied to human capital because it relies on behavior-change processes.”

2. Ensuring validity and contextual fit of AI tools

Building frameworks for validation and evaluation

When it comes to scaling AI in public health, one of the most pressing gaps lies in how solutions are validated. Around the world, very few robust **frameworks exist that can both assure scientific rigor and meet the urgency of health system needs.** In practice, this often leaves health care providers adopting tools without clear markers of quality or safety. One promising model is the **use of challenge frameworks: open calls where solutions are tested transparently against defined use cases.** Such mechanisms create space for innovators to demonstrate value, while giving health systems confidence in the performance of tools before procurement. They can accelerate safe adoption by bridging the trust gap between developers and frontline practitioners.

In Catalonia, this idea has already been institutionalized. The regional health system has defined technical, ethical, and legal requirements for trustworthy AI, including validation metrics, bias mitigation, explainability, and compliance with medical device regulation. Each solution undergoes a formal evaluation process before deployment. Rather than banning the inevitable decentralized adoption of tools such as generative AI, Catalonia promotes “responsible autonomy” by issuing best-practice guidance, maintaining an AI observatory to track use, and investing in training for both professionals and citizens.

Taken together, these approaches highlight the need for frameworks that are both rigorous and adaptable, balancing top-down governance with bottom-up realities, and ensuring that AI adoption is safe, trusted, and sustainable.
Finally, although many countries with clear digital health or AI strategies demonstrate appetite for innovation, regulatory frameworks often struggle to keep pace. Without integrating evaluation and compliance considerations from the outset, solutions risk being delayed or undermined by legal and ethical challenges later in the implementation process.

Adapting AI through localization and contextual fit

For AI tools to succeed in public health, localization is the difference between adoption and abandonment. A system designed for one setting can collapse in another if local realities are ignored. Even small details, such as how geographic units are defined, can determine whether health workers feel the system belongs to them. Competing priorities further complicate adoption—a caregiver managing a child with a life-threatening illness, for instance, cannot be expected to pause and download an app, no matter how well-designed it may be. Trust is fragile: **many frontline staff recall past pilots that never scaled, leaving behind unused devices and eroded confidence.** To rebuild it, solutions must be intuitive, responsive to local workflows, and visibly add value from day one.

> “This is what AI can look like when you co-design: only three minutes of training, offline functionality, with outputs that are understood by non-experts.”

Localization is not merely a technical challenge, it is deeply political, cultural, and human. Funders and implementers increasingly stress that co-design with users like ministries, health workers, and even patients must happen early and often. This creates ownership and lowers resistance. Countries that align AI adoption with their broader health strategies, and where governments are visibly energized by the potential of AI, can move faster from pilot to scale. In short, localization means embedding solutions in the grain of the health system, not laying them over the surface.

3. Designing strategies that enable system-wide scale

Establishing conditions for system-wide adoption

Scaling AI in public health is not simply a matter of replicating a successful pilot. **Scale must be designed in from the very beginning,** built into the architecture of the solution, not bolted on afterward. Strong algorithms are necessary but insufficient; **what determines whether a system grows or stalls is how well it fits the realities of the health systems it serves.** One principle stands out above others: design for the hardest contexts first, not the easiest.

BOX 1 DESIGNING FOR SCALE FROM SIERRA LEONE

In Sierra Leone, where maternal mortality is among the highest in the world, developers began by studying how CHWs already screened pregnancies before introducing an AI-enabled ultrasound prototype. If existing methods were effective, AI would add no value. This mindset of questioning the need for technology before building it shaped a product that worked with the system rather than against it. The design prioritized usability: a traffic-light risk interface, only three minutes of training required, and offline capability for areas without electricity or connectivity. Scaling also meant thinking beyond the device with each development stage including go/no-go checkpoints, and; rather than building a siloed solution, the team digitized antenatal care data directly into national health systems through smartphones that were deployed by BabyChecker in the field.

The result demonstrates that scaling from the outset means designing for real-world conditions, intuitive use, and system integration such as solving problems at multiple levels—not just detecting risks but strengthening the care pathways and data infrastructure that underpin them.

Furthermore, contextual fit begins with efficiency. AI gains traction where it addresses clear bottlenecks, relieving documentation burdens, supporting overstretched workforces, or improving diagnostic accuracy in under-resourced settings. But **efficiency alone is not enough.** Adoption must be needs-driven and evidence-backed, validated against local datasets, and aligned with the regulatory and procurement systems that govern how health tools are approved and funded. Catalonia's shift from pilots to procurement illustrates this well: by anchoring AI deployment to identify gaps in primary care and ensuring compliance with local standards, the system built the institutional confidence necessary for sustained investment.

Financing is another condition that cannot be assumed. Tools like CAD4TB only achieved scale once their pricing fit within donor and government procurement frameworks. **Affordability and long-term sustainability are not secondary considerations, they are prerequisites.** Without them, even the most effective tools remain confined to project budgets and donor cycles, never becoming part of the permanent health system infrastructure.

At the frontline, adoption hinges on whether health workers see value from the very first interaction. When a tool saves time, simplifies reporting, or sharpens diagnostic accuracy, ownership follows naturally, and that ground-level ownership is what sustains adoption as systems grow. Embedding AI into existing platforms, such as electronic health records (EHRs) or national data systems, removes the friction that typically stalls uptake. Kenya's experience with AI decision support integrated directly into EHRs demonstrates how seamless workflow integration can transform skepticism into acceptability.

> "Iterative rollouts give us the chance to see what went wrong, what succeeded, and make changes quickly before waiting six months or a year."

Ultimately, while policymakers control budgets and institutional mandates, it is health workers and patients who determine whether a system scales or stagnates. Designing for scale from the outset means designing for the people closest to the point of care. Their trust, earned early and maintained consistently, is what transforms a promising pilot into a system that holds together under the pressures of national-level implementation.

Aligning donor perspectives with long-term integration

From a donor's perspective, assessing whether an AI solution is worth funding goes beyond evaluating the technology itself. What matters most is contextual fit and the credibility of a pathway to long-term integration. The first signal is **co-design:** solutions built alongside frontline workers, ministries, and patients from the outset arrive with built-in ownership, lower resistance, and far greater prospects for sustained adoption. But design is only the starting point. Funders also scrutinize the broader ecosystem, the **strength of partnerships** across government, academia, and civil society, and crucially, whether political will exists to transition a donor-funded pilot into a program that a health system will own, finance, and maintain long after external support ends.

> "Funding agencies are critical to sustainable scale-up ... sustainability has to be built in from the start, so governments aren't left exposed when funders step out."

Sustainability serves as another critical filter. Donors look for evidence that governance and financing models have been considered from the outset, not assembled hastily as a project approaches its end. **Alignment with national strategies** is equally decisive: solutions that map onto government priorities are far more likely to survive the withdrawal of donor funding and transition into permanent system infrastructure. Where governments are actively energized by AI's potential, as seen in Rwanda and Kenya, the conditions for scale are demonstrably stronger, and donor investment carries a higher prospect of lasting impact.

Ultimately, for donors, the value of an AI solution lies less in its novelty and more in its **ability to embed, endure, and evolve within the health system it serves.** The most fundable tools are not necessarily the most sophisticated, they are the ones most likely to still be running, and still improving care, long after the grant has closed.

Fostering a culture of scale and sustainability

Bureaucratic silos within ministries, where programs such as TB, HIV/AIDS, and malaria operate independently, creating fragmentation that hinders AI adoption. Without unified overall sight or clear national AI policies for health, each department pursues isolated initiatives, making systemwide integration nearly impossible. **Scaling AI in public health requires cultivating trust, ownership, and the right enabling environment.** On the ground, health workers often resist new tools, fearing job loss or disruption to familiar workflows. Digital literacy varies sharply, especially beyond major cities, and trust is fragile after years of pilots that failed to stick.

In summary: building a culture of scale

Scan here to watch the full webinar:

To build a culture of scale, engagement must begin early. Tools should be tested directly with end users through acceptance testing, ensuring they fit existing workflows rather than forcing disruptive change. Equally important is showing immediate value: reducing time spent on tasks, improving diagnostics, or simplifying reporting. When health workers see tangible benefits, they begin to take ownership.

> "Prioritizing real system needs is essential ... the greatest value of AI often lies in redesigning workflows, improving coordination, optimizing processes."

Policy and infrastructure are equally critical to the success of AI integration in health systems. Governments must invest in strengthening digital infrastructure, developing coherent policy frameworks, and dismantling programmatic silos that hinder coordination. Training initiatives should be responsive to varying levels of digital literacy, ensuring that tools are intuitive, accessible, and readily adopted by frontline workers. Implementation processes are most effective when iterative, beginning with small-scale rollouts, followed by rapid feedback, adaptation, and gradual expansion. Such an approach fosters both confidence and momentum, creating the institutional and cultural conditions necessary for sustained adoption. Ultimately, a culture of scale emerges when both systems and individuals are empowered, supported, and persuaded of the tangible value that AI contributes to health care delivery.

Webinar 3

Financing the AI ecosystem for health

Co-authors: Ross Hawkins | David Lawson | Alex Grant | Robert Cryer | Eric Sutherland | Resham Sethi | Deveshi Roy

Financing AI for health requires more than funding isolated tools. To move from pilots to full integration, governments and donors must align clear reimbursement and procurement pathways, adopt phased financing models that reduce political risk, and embed safeguards to ensure rural and underserved populations benefit alongside urban centers. Regulatory clarity must be treated as a lever for reducing risk rather than an obstacle, investments in capacity building and governance remain essential to ensure that tools are used effectively.

Key discussion points:

- **Government leadership and policy clarity** through clear financing roles and embedding AI into core budgets.
- **The role of private and philanthropic partners** beyond capital, including market shaping and procurement levers.
- **Financing models and evaluation approaches.**
- **De-risking AI investments** through regulatory clarity, stakeholder engagement, and phased co-investment.
- **Operational readiness** and hidden costs, including systems, people, processes, and the real total cost of ownership.

1. Government leadership and policy clarity

Whole-of-government ecosystem planning

A coordinated, whole-of-ecosystem approach enables public and private actors to adopt, adapt, and scale AI in alignment with national health priorities and digital standards. A central coordinating mechanism or national digital blueprint is essential to align investments across federal, state, and local levels, support interoperability, and enable continuous learning and improvement. Effective adoption requires coordinated investment across multiple layers of government, not just in siloed tools, but also in connectivity, interoperable data systems, skilled workforce capacity, and long-term operational structures.

BOX 1 AUSTRALIA'S WHOLE-OF-GOVERNMENT DIGITAL AND AI LEADERSHIP

Australia's Digital Transformation Agency leads with its Digital Investment Plan, guiding government agencies to plan and implement digital investments effectively. Complementing this, Australia's GovAI initiative drives a whole-of-government approach to adopting AI safely, securely, efficiently, and ethically across public service agencies. **Together, these efforts show how national leadership can foster collaboration, build capability, align investments, and accelerate progress by maximizing return on investment while upholding public trust through clear guardrails.**

National vs local financing roles

Clarity in financing responsibilities emerged as a core issue. At the national level, governments are expected to fund shared public goods: digital public infrastructure, standards, national procurement frameworks, and priority platforms. These create consistent pathways for suppliers and give confidence to investors. Local systems, in contrast, remain the primary buyers and implementers, tailoring solutions to their contexts. To support this division of roles, **reimbursement clarity is critical.** In practice, this means establishing rules-based pathways and standardized technology assessments that ensure predictable funding and adoption routes. When assessment bodies recommend technology, financing and implementation are more likely to follow, creating greater certainty for suppliers and decision-makers.

Embedding AI into core budgets

Governments are often reluctant to finance large IT projects, which makes it difficult to justify AI investments unless framed within broader productivity and economic policy goals. One approach is to **quantify the cost of inaction,** which includes productivity losses, health system inefficiencies, and societal costs of not adopting effective AI solutions. Building the business case around these avoided losses strengthens arguments for initial and sustained financing.

> “By calculating that cost of inaction, there's benefit in actually taking time to unpack the specific workflow”

A phased approach to budgeting was also emphasized. Rather than large upfront allocations, staged investments tied to measurable milestones allow governments to clearly define success and pace implementation, test and quantify outcomes through a shared evaluation framework, build confidence, and integrate AI progressively into core health budgets.

BOX 2 JUSTIFYING AI INVESTMENT THROUGH THE COST OF INACTION IN AUSTRALIA

A recent feasibility study by Australia's Therapeutic Goods Administration (TGA) found that **adopting AI and automation would enable the agency to keep pace with the growing volume and complexity of applications submitted by medicine sponsors.** This is expected to improve staff well-being and productivity and allow evaluators to draw on historical and current intelligence to strengthen medicine regulatory processes. The business case was further strengthened by benchmarking against domestic and international AI case studies, making explicit the opportunity costs and risks of inaction.

2. Role of private and philanthropic partners

Governments and philanthropic organizations both influence the shape of AI markets. Several levers stand out.

Predictable reimbursement and regulatory clarity give private actors the confidence to invest in solutions that can scale across health systems.

Value-based procurement shifts purchasing decisions toward tools that demonstrate measurable outcomes rather than just technical novelty.

Shared infrastructure and buyer guides seeded by government or philanthropic actors reduce search costs, increase transparency, and help stakeholders identify solutions suited to local needs.

Together, these mechanisms create environments where financing flows more easily, solutions scale more predictably, and health systems are less likely to be left with fragmented or duplicative tools.

Private and philanthropic actors also have roles that extend well beyond providing capital. Their contributions can accelerate adoption when they bring technical expertise, support knowledge transfer, and create safe environments for testing new approaches. Equally important is their ability to strengthen governance by funding standards development, offering technical advisory support, or convening diverse stakeholders around best practices. But the most durable impact comes from local capacity-building. When partnerships invest in training health workers, transferring technology, or embedding skills in-country, they leave behind systems that can sustain and expand AI use long after the initial funding cycle has ended.

3. Financing models and evaluation approaches

Matching finance instruments to tech maturity

Different financing instruments suit different stages of AI maturity.

Grant-based funding plays a role in early-stage research and pilot projects, especially where equity and innovation are priorities but scalability has not yet been proven. For infrastructure-heavy initiatives, such as platforms that require data integration or long-term operational capacity, blended finance models were highlighted to de-risk investment. These combine public and private contributions, reflecting the fact that **public funds generate private benefits and vice versa.**

At the most mature stage, outcome-based contracts were described as the gold standard, aligning incentives across governments, providers, and investors. Such models are best suited for long-term, population-level interventions like chronic disease management or preventive health programs. However, this requires robust data systems and clear metrics over extended periods to uphold accountability and effectiveness.

Proportionate economic evaluation

Countries often face challenges in building full economic cases for AI because primary data is incomplete or unavailable. In these contexts, **pragmatic approaches were emphasized.** Policymakers can use secondary data, comparable proxies, or international benchmarks to approximate return on investment. Even without perfect baselines, this allows them to produce indicative estimates of value and avoid decision-paralysis by lack of evidence.

Workflow mapping was identified as a practical tool for capturing tangible opportunities and benefits by identifying the pain points, current technologies used, processes and people involved at each stage. By analyzing how AI affects specific processes such as reducing “touch-time” for staff or shortening “cycle-time” for administrative or diagnostic tasks, business cases can demonstrate productivity gains in concrete terms.

BOX 3 TIPS FOR PROPORTIONATE ECONOMIC EVALUATION

- Consider how AI will support the design of evidence-based policy, improve workforce productivity and planning, and build system readiness to prepare for the future.
- Evaluate the impact on citizens’ health outcomes and system costs (e.g., the potential for AI to reduce hospital and residential aged-care admissions, shorten stays, reduce errors, and lower acute and primary health care expenditure by streamlining processes and keeping people healthier for longer.
- Apply systems thinking and analysis at both the micro and macro level equips decision-makers to weigh short-term costs against long-term efficiency and societal outcomes, including quality-adjusted life years and broader public health, social, and economic gains.
- Incorporate second-order impacts of AI on the social determinants of health is where multi-factor productivity arguments start to gather momentum.

Phasing pilots was seen as another way to ensure proportionate evaluation. Rather than committing to large-scale investment upfront, countries can prioritize smaller, high-utility, low-effort pilots first, using them to generate evidence that informs subsequent funding decisions. The focus should be on scalability. Pilots that demonstrate novelty but cannot integrate into national strategies add little value in the long term. Funding must prioritize solutions designed to expand across geographies and service levels, strengthening the system (also a theme from Webinar 2).

Equity and access imperatives

Poorly structured AI financing risks deepening geographic and social inequalities, with urban hospitals and private providers capturing the gains while rural and public systems and culturally diverse groups left behind. National investment in digital public infrastructure such as health information exchanges, standardized data protocols, and shared connectivity layers are essential to prevent this. These foundations create a level playing field so that AI tools can plug into the wider system, rather than operating as isolated, urban-only pilots.

BOX 4 MANAGING BIAS AND INEQUITY RISK THROUGH RESEARCH

Targeted inclusion measures

Equity also requires deliberate action to extend AI benefits to underserved populations. Several mechanisms stand out:

- **Targeted subsidies or funding** for rural and remote facilities to ensure AI services reach beyond well-resourced hospitals.
- **Equity checks in procurement** that embed criteria favoring tools designed for low-resource contexts rather than those suited only to advanced systems.
- **Buyer guides and evidence frameworks** that compare available solutions, highlighting evidence-based options, and reduce the risk of inappropriate purchases.
- **Funding research and partnerships with community-controlled organizations** to assist culturally and linguistically diverse organizations in guiding, implementing, and supporting the culturally safe use of AI and broader digital health technologies.

Action on equity

To better respond to cultural diversity and manage risks of bias and inequity, **Australia invested in building an evidence base to understand the applicability, concerns and opportunities of AI to Aboriginal and Torres Strait Islander health care.** This research establishes Aboriginal and Torres Strait Islander organizations as self-determined leaders in this space and offers discussion of what a culturally self-determined AI health ecosystem can look like.

De-risking AI investments

Regulatory clarity as a de-risking lever

Regulation, often seen as an obstacle, can instead be a tool to lower investment risk. Fast-track pathways that are proportionate to the level of risk allow innovative AI tools to enter the market without unnecessary delay. The safety net comes from strong post-market surveillance, ensuring technologies continue to perform as intended once deployed.

Equally important is coordination across jurisdictions. **Mutual recognition between trusted regulators avoids duplication, reduces compliance costs, and shortens the time needed for technologies to reach multiple markets.** Regulatory systems that are predictable, proportionate, and collaborative give both governments and private investors the confidence to finance AI adoption at scale.

> "Regulation should be seen as a positive, because we regulate to ensure patient safety."

Engaging with early adopters, end users, and diverse groups

Engaging with early adopters of technology both domestically and internationally can help de-risk proposals by providing tangible benchmarks for expected investment, prerequisites, and benefits. This approach is particularly **valuable for funders who are hesitant to be first movers, as it allows them to learn from others' experiences and adopt proven practices with greater confidence.**

Actively consulting on challenges and opportunities, and co-designing solutions with system and **end users** with lived experience is essential to reducing risks and unlocking the full benefits of AI. **This inclusive approach embeds both the health equity and ethical lens from the outset,** helping not only to prevent unintended consequences but to actively close existing gaps.

When preparing funding proposals, it is valuable to capture **stakeholders' perceptions** and proactively address potential risks by clearly explaining how, where, and why AI will be implemented in alignment with national safeguards. Ensuring human oversight and validation is essential for maintaining both internal and public trust. When concerns arise about AI replacing human decision making, it is important to clarify that AI is being used as a supportive tool that enhances people's roles, boosts productivity, improves well-being, and strengthens system responsiveness.

Phased financing and co-investment

Financial structuring was highlighted as another way to share and manage risk. Key approaches include:

Phased financing tied to milestones, so investments scale only after evidence of value is demonstrated.

Co-investment pools that combine public, donor, and private funds, spreading risk while aligning interests.

Pooled procurements and shared testbeds (including joint international investments) that allow multiple actors to validate solutions collectively before committing to full roll-out, and benefit from accelerated learning.

These approaches reduce the likelihood of governments shouldering the burden of unproven technologies on their own, while still creating space for innovation to grow.

BOX 5 AUSTRALIA'S MEDICAL SCIENCE CO-INVESTMENT PLAN

The Australian government has delivered a Medical Science Co-investment Plan to outline opportunities for government and industry to leverage Australia's strengths and target areas with high economic potential.

Developed with industry, this plan identifies digital health, including AI and ML, as one of the key investment opportunities to strengthen Australia's medical manufacturing capabilities.

Strengthening LMIC proposals and fit

Financing decisions often hinge on how well proposals align with national strategies and health system realities. **Strong submissions show clear integration into national digital health strategies and enable the delivery of government's strategic health priorities such as prevention and early intervention and integrated health systems.** This demonstrates that new tools will not remain as isolated, one-off pilots but become part of broader system architecture and a key strategic asset to deliver on short- and long-term national health priorities. Compliance with recognized standards, such as cyber security and interoperability frameworks, is another marker of readiness, **reducing risk for funders who need assurance that tools are safe and can be scaled.**

Funders also look for evidence of sustainability. This includes capacity plans that account for workforce training, governance structures, and long-term operations. Proposals that achieve economies of scale and highlight how AI can plug into existing health information exchanges and national roadmaps signal that investments will not vanish after initial funding cycles. In short, system fit and long-term viability weigh heavily in financing decisions.

4. Operational readiness and hidden costs

Working with and around legacy systems

Countries face very different starting points in adopting AI. Those without legacy systems—old EHRs, fragmented databases, and bespoke hospital IT—may appear better positioned to leapfrog, but they still require essential infrastructure: stable electricity, reliable connectivity, and the ability to generate and manage quality data. Without these basics, even advanced AI tools remain unusable.

For countries with entrenched legacy systems, the challenge is different. Here, investment must focus on interoperable tools and phased migration strategies. Rather than replacing entire architectures, AI solutions should be designed to sit atop existing systems while gradual upgrades unfold. This approach allows governments to modernize incrementally, balancing the need to harness AI's potential with the reality of long-standing technical debt.

Real total cost of ownership

AI feasibility studies often underestimate the resources required to sustain adoption. While technology licenses or devices are visible costs, **the real expenses lie in the supporting ecosystem.** These include:

1. Data cleaning and ongoing quality management
2. Process redesign and workflow adjustments
3. Workforce training and re-skilling
4. Continuous validation and algorithm monitoring
5. Governance and compliance oversight
6. Long-term system maintenance

If such costs are ignored, projects risk failing after initial pilots. Modelling should also capture the iterative nature of spending: early pilots, scale-up phases, post market surveillance, and continuous improvement cycles each draw resources. By making these costs explicit, feasibility studies provide a more realistic basis for investment decisions.

People, process, and change management

Technology rarely fails because of code; it fails because people and processes are unprepared to use it. Resistance to new workflows, role confusion, and the absence of local champions were all identified as primary reasons why AI tools stall after pilots.

> "Around 70% of AI adoption, it is not a technological issue. It is a people and a process issue."

Operational readiness therefore depends on parallel investment in organizational change. Training programs, end-user piloting, and the cultivation of internal advocates are critical to build trust and ownership. Investment in process and role redesign is essential to enable AI to meaningfully augment business as usual to realize its full benefits. In doing so, AI can be embedded safely by maintaining appropriate human oversight and final decision making. This has the potential to significantly improve productivity and staff well-being by removing mundane, repetitive, and time-consuming tasks and recasting roles to focus on using the skills and knowledge the staff was hired for in the first place.

Supporting staff to integrate basic AI into their everyday work is a key first step to bringing people along on the journey and building the buy-in required for a broader roll-out. This may require investing in current systems and infrastructure first to demonstrate the value AI brings, while building capability to balance literacy and trust with vigilance.

In summary: financing for scale

Financing AI for health requires more than funding isolated tools. To move from pilots to full integration, governments and donors must align clear reimbursement and procurement pathways, adopt phased financing models that reduce political risk, and embed safeguards to ensure rural and underserved populations benefit alongside urban centers. Regulatory clarity must be seen as a lever for reducing risk, not as an obstacle, while capacity building and governance investments remain essential to ensure tools are used effectively.

Scan here to watch the full webinar:

Webinar 4

AI and the invisible workforce: task shifting and automation risks

Co-authors: Anima Rawat | Kehinde A. Muraina | Katarina Vujovi | Manoj Sanker | Ousseni Ouedraogo | Sharan Murali Sarah Parwez Eric Sutherland | Resham Sethi | Pallavi Jain| Deveshi Roy

As AI continues to evolve, the intersection between AI and the health care workforce broadens. This interaction is transforming how care is delivered by redefining roles, redistributing tasks, and shifting the boundaries between human and technology collaboration.

While AI-integrated health systems hold great promise of relieving overstretched systems, their rapid adoption also carries the risk of deepening existing inequities if integration is not carefully managed. Technology evolves faster than systems can keep up, and many countries don't have optimized health data. Global health systems are already operating under strain, with the WHO projecting a shortage of 11 million health workers by 2030, predominantly in LMICs. Against this backdrop, AI offers a timely opportunity to augment human capability but only if it is designed and deployed with empathy, trust, transparency, inclusivity, and iterative design.

> "Artificial intelligence is not there to replace health workers but it's there to support them."

Key discussion points:

Opportunities for strengthening the health workforce through automation of administrative burden and task shifting.

→ **Perception of health care workers** toward integration of AI tools.

→ **Managing accountability and ethical risks** across the AI development and deployment lifecycle and recognizing the invisible labor behind AI systems.

1. Opportunities for strengthening the health workforce

AI's most immediate impact lies in alleviating the chronic workload pressures facing health care workers. In many countries, administrative duties consume a colossal portion of health care workers' time, detracting from patient interaction. AI-powered automation, such as digital scribes and clinical note summarizers, can reduce these inefficiencies dramatically and can help take on more than 30 percent of time-consuming tasks.

BOX 1	AI SCRIBES TO ALLEVIATE THE ADMINISTRATIVE BURDEN

In the last year, there has been a 600 percent increase in AI scribe adoption, helping save 1 minute per visit. Pilot studies have demonstrated that AI-assisted scribes in US hospitals have saved more than 15,000 clinician hours annually, improving both productivity and morale.

Additionally, AI tools can help optimize the workflow of workers, stratify risk, and augment their skillsets. Automation, when used strategically, offers a chance to reimagine health work by freeing up time for caregiving, problem-solving, and patient engagement while allowing AI tools to handle repetitive and routine tasks. However, to truly harness the opportunities AI presents, countries must make deliberate investments in training health care workers and building appropriate tools and accountability mechanisms to protect against the risk of AI tools exacerbating burden of the workforce.

2. Perception of health care workers toward integration of AI tools

Anxiety, adoption, and the human factor

In spite of the considerable potential benefits and efficiency gains that AI tools promise for health care delivery, anxieties amongst health care workers around AI integration continue to remain high and lead to adoption resistance. Many fear that the introduction of AI could compromise the quality of care, displace clinical judgement, or render their roles redundant.

A key structural challenge contributing to this resistance is the fragmented nature of health programs, which frequently operate in silos with limited interoperability. As a result, frontline healthcare workers are often required to navigate multiple digital portals simultaneously for data reporting, beneficiary tracking, and program monitoring, often in settings with limited internet connectivity. Workers have pointed to the burden of managing "too many portals to update" as one of the most significant barriers to the sustained adoption of digital and AI tools.

Compounding this, many of these **tools are developed without sufficient attention to local context or the day-to-day workflows of the workers expected to use them,** creating friction that deters both usability and uptake. Workers have further highlighted how recurring challenges, including frequent and disruptive software updates, defunct or non-functional portals, and shifts in political commitment and program priorities, have progressively eroded trust in these tools over time.

BOX 2	OECD EVIDENCE ON TASK SHIFTING AND EMERGING HYBRID ROLES

Empirical evidence from OECD countries shows that AI adoption typically leads to task shifting rather than job displacement. Roles evolve through positive augmentation of existing workforce capabilities rather than disappearing. New hybrid roles that blend clinical expertise with technology practice have emerged (e.g., chief medical information officers, digital platform managers, digital nurse specialists, and AI-safety officers. OECD projections indicate that demand for such hybrid roles could rise by 40 percent by 2030, reflecting the need for human oversight even in highly automated environments.

Trust grows when AI tools are seen as stable, co-created, and directly responsive to frontline realities. Workers need to understand the "why" behind automation. This helps them to be far more likely to embrace it as an ally rather than resisting it.

> "We can save 20 percent of a nurse's time simply by taking away the registers and digitalizing their workflows, but trust breaks when tools change too often."

Building digital skills and confidence

Reskilling the workforce requires a structured, sustained approach, such as remodeling it as a part of refresher training. Similarly, health care workers need to be **willing and flexible to adapt to new technologies.** Many frontline workers still rely on personal devices and self-funded data plans, limiting access to digital systems. Embedding device procurement, maintenance, and localized technical support within budgets is essential to ensure fair participation. Upskilling should also not just entail endpoint solutions but rather **make health care workers understand what the technology is and how it can help them.** Digital literacy remains a foundational layer before any AI skill-building can be facilitated for LMICs.

In addition to building digital literacy, **health care workers must also be equipped with the confidence to use AI tools.** Many frontline and CHWs, particularly those from rural and low-connectivity regions, remain apprehensive about using smartphones and digital platforms. Addressing this apprehension through targeted training and the cultivation of soft skills is as vital as the technical infrastructure itself, and ultimately what determines the success of AI integration and guards against rejection.

> "Lead with empathy. Design for the end user who will actually use your solution to help reduce the inequity that comes with sacrosanct structures that we function with."

3. Mitigating adoption resistance through co-creation

Designing solutions with health care workers to meet contextual requirements

Field experiences illustrate how context can determine the success or failure of AI deployment. Apprehensions around being replaced or tools being too complex and not reflecting community needs often deter health care workers in the field. Solutions designed in controlled environments often falter when introduced into busy, under-resourced clinics where conditions bear little resemblance to stimulatory conditions. In contrast, solutions that are co-designed with end users, contextualized to local needs, and refined through iterative testing consistently achieve greater uptake and sustained use.

BOX 3 CO-DESIGN AND LOCALIZATION IN PRACTICE: LESSONS FROM IVORY COAST AND INDIA

Field experiences from the Ivory Coast recommended solutions that incorporate features such as offline functionality or audio-voice in the interface. Local languages can increase both the confidence of the healthcare worker and the patient in their responses.

In India, AI-enabled newborn monitoring applications such as Shishu Maapan allow community workers to assess infant growth using short videos, eliminating the need for physical weighing scales. Similarly, smartphone-based cough analysis tools for TB screening are helping expand screening coverage in areas lacking X-ray infrastructure. In field trials, these systems increased case detection significantly, showing how AI can expand the reach of limited diagnostic resources.

Studies conducted by the OECD that analyzed the usage and impact of AI tools on health professionals reveal that the “human by design” factor is often missing within AI tools. Introducing real-time feedback mechanisms by health care workers, such as the case of Estonia’s digital sandboxes, can minimize gaps between AI technology and clinical reality. In India, the SIMPLE application for the India Hypertension Control Program, co-designed entirely with nurses, successfully onboarded 2 million patients. Actively engaging field staff in the ideation and co-designing of AI tools can help build solutions that support their workflows and day-to-day realities.

Preserving human-centered care

As AI tools take on a greater share of diagnostic and decision-support tasks, it becomes imperative to preserve the human connection that underpins effective and holistic healthcare delivery. AI tools should never compromise clinical judgment or impede with the caregiving process.

AI innovators must be careful not to diminish the “human in the loop” and conduct usability, acceptance and feasibility studies with public health institutes and end users to flag friction points. Tools should augment health care workers’ decision-making capacity rather than making the decisions for them. Design must also account for the broader care architecture, recognizing that solutions built for one type of health care worker will interact with, and affect, others across the system.

Research from high-income countries indicate that patients are most comfortable when doctors use AI tools as assistants as opposed to autonomous decision-makers. The critical opportunity here is for the AI tools to strategically fill a void where a specialist is not available by extending diagnostic and clinical capabilities to a different cadre of health care providers. For example, an AI screening tool for diabetic retinopathy was deployed in India due to a lack of trained ophthalmologists.

While AI models that surpass the capabilities of human specialists do exist, the underlying potential risk of overreliance must not be overlooked. While these tools can offer greater accuracy, interpretability often diminishes.

4. Managing accountability and ethical risks

Accountability though governance frameworks

AI’s increasing role in clinical and administrative workflows introduces new layers of ethical and operational complexity. Errors generated by algorithms, such as incorrect clinical summaries or biased diagnostic predictions, can have severe consequences.

Ensuring accountability requires clear governance frameworks and shared responsibility across the health ecosystem. Model developers, implementers, and end users all hold pieces of the accountability chain. Nigeria’s Design Health Initiative mandates guidance and a governance structure framework for the use of digital solutions in the country.

Risk mitigation must include transparency in model source, local validation of data, and clearly defined lines of responsibility.

> “Everyone should be held accountable both from top to bottom—the governance, the government, regulators, health care workers, and even the makers of this solution. Everyone should play critical role in ensuring that the AI solution is ethical and maintains the highest quality.”

Human oversight

Human oversight remains non-negotiable in high stakes clinical environments and is fundamentally a **workforce strategy.** As AI systems take on greater roles in diagnosis, treatment recommendations, and administrative decision-making, **the presence of qualified human judgment at critical decision points is essential to patient safety and ethical accountability.**

Oversight mechanisms should be embedded across the entire AI deployment lifecycle, from design and validation to monitoring and review. Clinicians and administrators must retain the authority to question, override, or escalate AI-generated outputs, particularly in cases involving vulnerable populations or irreversible decisions.

Effective human oversight also requires adequate training. End users must understand the limitations of AI tools, including the risk of automation bias and the tendency to defer uncritically to algorithmic outputs. Governance frameworks should define clear protocols for when human review is mandatory and establish feedback loops that allow oversight findings to inform model improvement over time.

Recognizing the "invisible workforce"

Behind every functional AI system lies a vast, often invisible workforce of data labelers who clean, filter, and tag the data that trains AI models, yet their contribution is rarely acknowledged.

Recognizing this digital labor, as well as emerging roles such as hybrid and back-office workers, is central to ethical AI development. Integrating these workers more into existing workflows begins with giving them visibility and including them in data collection processes that can inform how they are upskilled and supported. Countries are now beginning to integrate digital labor into workforce planning, acknowledging it as an integral part of the health technology value chain.

> "The first step is to make them visible in the language we use."

In summary: the workforce comes first

Scan here to watch the full webinar:

AI's most immediate contribution to the health workforce lies in reducing administrative burden, freeing clinicians time for engaging in patient care. OECD evidence consistently shows that AI shifts tasks rather than eliminating jobs, with demand for hybrid clinical-technology roles projected to rise sharply by 2030. For this potential to be leveraged, co-designing with frontline workers is non-negotiable as these tools must fit clinical workflows and local realities in . Addressing adoption resistance requires showing immediate value, building digital confidence alongside literacy, and maintaining human oversight in all clinical decision pathways. The invisible workforce which included data labelers and back-office AI workers must be recognized, planned for, and supported as an integral part of the health AI ecosystem.

Webinar 5

Regulation and ethics: from guidelines to enforcement

Co-authors: Rohit Malpani | Vidushi Chaturvedi | Dr. Osama Elhassan
Paul Campbell | Anusha Purushotham | Eric Sutherland
Aishwarya Rohatgi | Abhinav Goyal | Deveshi Roy

AI is entering health systems at a speed that existing governance structures cannot match. Principles on safety, fairness, and accountability are well understood, yet most systems still lack the tools to enforce them. This gap between ethical intention and operational capability is now one of the most significant challenges in responsible adoption of AI in health.

The insights discussed highlight recurring pressures: rapid innovation, uneven regulatory maturity, limited digital infrastructure, and growing expectations from both implementers and users. These pressures expose structural weaknesses: research oversight mechanisms don't map cleanly to AI, data protection rules are inconsistently applied, and accountability is often unclear across the AI lifecycle.

This webinar session examined the major themes that emerged from these insights:

The ethics-governance gap and the urgent need to translate principles into enforceable operational mechanisms.

Operationalizing accountability through enforcement capacity, continuous feedback loops, and structured testing environments.

Regulatory asymmetry and capacity gaps, including the mismatch between innovation pressure and institutional readiness.

Data inequality and embedded bias, and the need for inclusive and representative validation.

Innovation pressure versus patient protection, and how to balance progress with safeguards.

Multi-stakeholder governance and collective accountability across the AI lifecycle.

Together, they outline what it will take to move from high-level principles to practical, enforceable systems that uphold safety, trust, and equity in the use of AI for health.

1. The ethics–governance gap

Ethical frameworks outpace real-world capacity

Ethical principles for AI in health are now widely articulated, yet the systems meant to apply them often lack the capacity to interpret or enforce them. Many decision-making bodies do not have up-to-date standards for AI, nor the technical expertise to evaluate complex models, shifting algorithms, or new forms of data use. **This leaves a structural gap between what is expected ethically and what is feasible operationally.**

This misalignment becomes most visible at early stages of the AI lifecycle, where research oversight is uneven, evaluation processes are outdated or not in alignment with present contexts, and system-level checks are built for static technologies rather than adaptive ones. The consequence of this misalignment is that the implementation of ethical guidance remains inconsistent.

Ethics in AI is still frequently treated as a conceptual layer rather than a technical or institutional requirement. In practice, ethical principles demand operational mechanisms: evidence standards, model auditability, transparent disclosure, and regulatory pathways that can adapt as technologies evolve.

> “Converting principles into enforceable expectations is now a core governance priority.”

The emerging insight is that ethics must be embedded early and repeatedly across the AI lifecycle. It cannot be retrofitted once a system reaches deployment, nor can it rely solely on voluntary compliance.

2. Operationalizing accountability

A recurrent tension emerges around where accountability should be placed. Developers often assume that responsibility lies downstream with health systems and regulators, while implementers argue that ethical safeguards must be built upstream into the design and development process. In contrast, it is also the case that excessive protection can stall innovation if there is no clarity on who the accountable party is. **This ambiguity allows risks to fall between institutional boundaries, creating blind spots in oversight.**

Clearer delineation of responsibility is essential, particularly for systems that can evolve over time or behave unpredictably. Without explicit requirements for transparency, performance claims, and risk disclosure, accountability defaults to systems not equipped to manage it.

Laws and guidelines lack enforcement strength

Across settings, AI in health is governed primarily through guidelines, advisory documents, and soft law rather than enforceable mandates. This limits the system’s ability to deter harmful practices, manage high-risk deployments, or prevent unsafe tools from entering the market. When compliance is voluntary, accountability becomes dependent on goodwill rather than oversight. Strengthening enforcement capacity requires both legal mechanisms and institutional infrastructure.

> “Without bodies empowered to monitor, verify, and respond to noncompliance, even comprehensive guidelines fail to deliver meaningful protection or assurance.”

AI tools deployed in health systems need continuous feedback loops

Accountability in AI cannot rely solely on pre-deployment evaluation; it requires ongoing monitoring. In the absence of oversight, it becomes increasingly difficult to evaluate impact (or lack of thereof) of AI tools, and without measurement, there can be no management. Traditional regulatory models do not accommodate systems that learn, drift, or degrade performance over time. Health systems therefore need structured feedback loops like real-time monitoring, post-deployment evaluation, and mechanisms to detect emerging risks. Such loops allow systems to identify when tools are underperforming, behaving unpredictably, or producing inequitable outcomes. **Without them, risks remain invisible until harm has already occurred, undermining both safety and trust.**

Accountability breaks down when responsibility is distributed loosely across actors without clear definitions. AI governance works only when roles are explicitly assigned, such as who evaluates risk, who validates performance, who monitors drift, who communicates limitations, and who intervenes when failures occur. Fragmentation leads to gaps that undermine safety, equity, and public trust.

Strengthening governance requires coordinated systems where responsibilities are aligned rather than layered haphazardly. Shared accountability must be designed, not assumed.

Controlled testing environments, including sandbox approaches and staged evaluation pathways, offer a mechanism to balance innovation with patient protection. **By allowing experimentation under defined constraints, these environments help assess safety, performance, and contextual fit before wide-scale deployment.**

This approach reduces system-level risk and enables clearer decision making about when and under what conditions AI tools should be integrated into routine care. It also provides developers with clarity on expectations, evidence requirements, and acceptable risk thresholds, helping align innovation with responsible practice.

3. Regulatory asymmetry and capacity gaps

Global frameworks require local adaptation

Most existing standards and governance models were developed in high-resource environments with mature regulatory institutions. Applying these frameworks in settings with different legal traditions, infrastructure constraints, or population needs can be both impractical and inequitable. **One-size-fits-all governance is ineffective when health systems vary widely in readiness, risk tolerance, and stakeholder expectations.**

As a result, countries increasingly require approaches that are context specific, allowing for the adaptation of ethical principles, performance requirements, and oversight mechanisms to local realities. This includes recognizing differences in data availability, digital maturity, linguistic diversity, and geographical variance.

Data protection, ownership, and sovereignty are foundational

AI in health depends on access to sensitive personal data. Many systems have begun defining patient privacy, ownership, and data localization requirements to protect against misuse and prevent data from flowing across borders without safeguards. These elements have become central to AI governance, shaping what can be built, where it can be deployed, and how risks are distributed.

> “Health data must be treated differently from all other data given that it represents both individual and collective interest.”

Clear rules around data sovereignty also form the basis for trust. When individuals and systems understand how data is stored, protected, and controlled, it becomes easier to evaluate AI tools and make informed decisions about adoption. Without this foundation, both innovation and public confidence remain fragile.

4. Data inequality and embedded bias

Global datasets systematically underrepresent large populations

Many widely used medical and imaging datasets disproportionately reflect demographic groups from high-resource regions. This bias carries through to training data, leading to models that perform well on certain populations but unreliably on others. When such systems are deployed in diverse or underserved contexts, they risk reproducing inequities rather than addressing them.

This imbalance is a governance challenge. It becomes crucial to trace and identify the source of bias in datasets and provide appropriate rectifications. Without deliberate strategies to diversify datasets, models may inherently disadvantage the populations most in need of improved health access.

Reliable AI in health requires validation across diverse patient groups, clinical settings, and geographies. Developers increasingly recognize that internal testing is insufficient and independent evaluation, external audits, and real-world performance studies are necessary to build confidence in model outputs. When AI is validated against clinically accepted reference standards and externally run assessments, its credibility strengthens substantially.

This approach also aligns ethical intent with technical execution. By embedding inclusion into training, testing, and validation, systems can reduce bias before deployment rather than relying on corrective measures afterward.

Even when AI models are technically robust, their benefits may not reach populations without sufficient digital infrastructure. Uneven availability of connectivity, devices, and clinical information systems determines who can access AI-enabled services and who remains excluded. This creates the risk that AI accelerates progress for some while widening disparities for others.

Addressing these gaps requires aligning AI deployment with broader investments in infrastructure, workforce support, and system strengthening. **Equity is not achieved through the model alone; it is achieved through the ecosystem in which the model operates.**

5. Innovation pressure versus patient protection

The speed of AI development exceeds the pace of oversight. New AI systems are being introduced faster than governance structures can interpret their mechanisms, assess their risks, or set appropriate safeguards, leaving health systems exposed to unexamined risks, especially when models behave unpredictably or evolve after deployment. The core concern is not innovation itself but the absence of synchronized oversight: without mechanisms to evaluate emerging tools in real time, health systems are forced into reactive positions, responding to harms or failures only after they surface.

Regulators in many settings face intense pressure to enable AI adoption, driven by expectations around efficiency, digital transformation, and economic opportunity, yet they often lack the expertise, tools, and staffing required to evaluate AI safely. The result is an asymmetry with rapid innovation on one side, limited regulatory capacity on the other, which, in turn, creates structural risk, as oversight bodies cannot meaningfully review claims, assess model performance, or anticipate system-level implications when proposals arrive faster than they can be analyzed.

Without deliberate investment in regulatory capacity, the gap between technological capability and institutional readiness will continue to widen.

> “The result is an asymmetry: rapid innovation on one side, limited regulatory capacity on the other.”

Commercial incentives do not always align with public health needs

Developers often prioritize rapid market entry and scale, which may conflict with the slow, evidence-based processes required for safe clinical use. As a result, tools can enter the health ecosystem without sufficient validation, transparency, or risk of disclosure. Health systems, meanwhile, bear responsibility for managing consequences they did not create and may not be equipped to handle. This dynamic can distort incentives. When innovation is judged by speed and market share rather than safety and equity, vulnerable populations face disproportionate risk. Aligning commercial ambition with public health protection is becoming an essential governance challenge.

Balancing innovation pathways with protective safeguards

Safe innovation requires controlled mechanisms that allow experimentation without exposing patients to avoidable harm. Structured pathways such as staged evaluation, monitored rollouts, and predefined performance thresholds help ensure that new technologies contribute to health goals while keeping risk within acceptable limits.

These safeguards serve two purposes: they protect populations, and they offer innovators clarity on expectations, reducing uncertainty around evidence requirements. When innovation is guided by predictable guardrails, both progress and protection become more achievable.

6. Multi-stakeholder governance and collective accountability

AI in health demands interdisciplinary collaboration

Designing, evaluating, and deploying AI in health requires insights from clinicians, technologists, ethicists, social scientists, community representatives, and policymakers. No single domain holds the full expertise needed to understand risks, interpret outputs, or anticipate the societal implications of these systems. Effective governance therefore depends on structured collaboration that integrates multiple perspectives from the outset. **Such collaboration ensures that tools are not optimized solely for technical performance but are shaped by clinical relevance, cultural context, and real-world feasibility.**

Public trust is a core infrastructure, not an outcome

Trust determines whether AI tools are adopted, used appropriately, and integrated sustainably into health systems. Trust is influenced by privacy protections, transparency about system behavior, clarity around limitations, and the lived experience of patients and providers. When these elements are weak, even highly capable technologies fail to gain traction. Building trust is therefore not only a communications task but a governance task. It requires consistent performance, demonstrable safety, and a clear commitment to equity. Without this foundation, AI cannot meaningfully contribute to improved health outcomes.

In summary: governing AI responsibly

Scan here to watch the full webinar:

AI in health is no longer defined by intention but by implementation. Ethical principles set out the direction, yet systems succeed or fail in the operational details: how evidence is generated, how risks are monitored, how responsibilities are shared, and how inclusive the underlying data and design processes are. The insights discussed point to a sector that recognizes both the transformative potential of AI and the regulatory, infrastructural, and institutional fragilities that determine whether those tools benefit or harm populations.
The future of responsible AI in health will rely on governance models that can adapt as quickly as technologies evolve. This requires investment in regulatory capacity, stronger feedback mechanisms, equitable data ecosystems, and collaborative decision making that spans technical, clinical, and societal expertise. AI can expand access, improve diagnostics, and reduce disparities, but only if systems embed accountability, prioritize equity, and treat trust as core infrastructure. The path forward is not only technical, but also institutional, ethical, and deeply tied to how health systems define their duty of care in an era of rapid digital change.

Webinar 6

Global health diplomacy and the AI frontier

Co-authors: Kavita Bhatia | Eric Sutherland | Mathilde Forslund | Bogi Eliasen Urvashi Aneja | Aishwarya Rohatgi | Abhinav Goyal | Deveshi Roy

AI is beginning to reshape how countries deliver care, build capability, coordinate public health responses, and negotiate shared standards. The session underscored that health systems are entering the AI era with uneven data foundations, variable digital maturity, and governance structures that were not designed for learning systems. Across contexts, foundational gaps persist: limited access to trustworthy datasets, lack of people-centric governance structures, inconsistent identifiers, fragmented records, weak interoperability, and care pathways that rarely match the assumptions built into most AI systems.

These constraints influence how safely and equitably AI can be deployed. In many settings, even assembling a representative dataset for local validation is difficult, clinical workflows vary widely, and frontline environments face capacity challenges that no model can compensate for. AI tools trained on data from high-income settings or other well-structured data risk misclassification, skewed performance, and reinforcement of inequities when deployed in systems with different operational realities. As several participants noted, without redesigning data infrastructure and governance alongside technology, AI may magnify rather than reduce structural gaps.

Building AI-enabled public health systems requires clarity of purpose, trusted institutions, and governance models that embed safety, rights, and accountability. Countries must strengthen their ability to evaluate, monitor, and stress-test AI systems in both regulatory sandboxes and real-world settings; align data practices with public-interest use; and develop legal, regulatory, and diplomatic mechanisms that support interoperability and protect populations.

As AI integrates across entire populations, cooperation across borders and convergence on shared guardrails will determine whether these technologies serve public value or deepen global imbalances.

> “AI assumes stable identifiers and interoperable systems—most health systems simply do not have that foundation.”

Key discussion points:

Data foundations and system readiness, including the structural constraints that predate AI itself.

Governance, safety, and regulatory capacity, including gaps in health-specific legislation and the need for continuous evaluation.

Purpose, public value, and system design by anchoring AI in long-term health goals rather than tool-centric experimentation.

Diplomacy, cooperation, and global alignment as structural requirements for safe and equitable AI on scale.

Trust, institutions, and public engagement as the foundation for responsible AI adoption.

1. Data foundations and system readiness

AI deployment depends on data systems that are **reliable, interoperable, and representative of real clinical and public health environments.** Most health systems, however, were digitized around facilities rather than data use, producing fragmented architectures that constrain responsible scale. Earlier digitalization efforts prioritized removing paper records within hospitals, leading to institution-specific technologies that do not align with national or population-level use cases. As countries transition into the AI era, these legacy structures shape whether models can be trained, validated, and deployed safely. Health records remain dispersed across institutions, structured differently, and difficult to reconcile. Facility-centric systems limit the ability to assemble longitudinal patient records or track care across levels, impeding the creation of **integrated datasets** required for AI-enabled diagnostics, prevention or public health surveillance.

Many AI systems assume **stable identifiers, consistent record-keeping, predictable referral pathways, and interoperability across providers.** These conditions rarely exist in low-resource or mixed-delivery environments. Health-seeking behavior is non-linear, pathways are fragmented, and record quality varies significantly—all conditions that affect how models learn and perform. Validating models on local datasets is difficult when hospitals or health care personnel lack statistical capacity to construct representative samples or curate clean validation sets. Even determining which cases to include or how to structure datasets becomes a barrier.

> “Our data problems are so deep that freeing data is still several steps away; we need to re-engineer the foundations.”

Without proper validation, AI systems risk misclassification, biased outputs, and inappropriate triage recommendations.

Equity, representation, and real-world risks to end users

AI's performance and value are fundamentally shaped by the environments in which systems are deployed. When data, workflows, and population characteristics diverge from the assumptions embedded in AI models, risks escalate. In many contexts, fragmented care pathways, limited data availability and structural inequities mean that **AI can misclassify, mistarget, or misrepresent populations, reinforcing rather than reducing disparities.**

Tools deployed without contextual calibration such as chatbots or decision aids, can create over-reliance without providing safe fallback mechanisms. This transfers risk onto frontline workers who may lack institutional protection, training authority to correct system outputs. In low-resource settings, weak oversight can amplify the consequences.

> “Without institutional capacity, risks fall on frontline workers and patients—exactly the people least equipped to manage them.”

Automation, surveillance and data governance

Automated data collection systems require strong governance frameworks to prevent unintended surveillance and erosion of public trust. Systems must balance improved data availability with privacy, control, and transparency, ensuring that data flows serve public health purposes without disempowering communities.

AI-enabled health futures will rely increasingly on real-time data captured through devices, sensors and patient-generated inputs. This shift demands architectures that can ingest, process and contextualize dynamic data while maintaining safety, accountability, and equity, especially in offline or low-connectivity environments.

2. Governance, Safety and Regulatory Capacity

Governance structures determine how effectively AI can be deployed, monitored, and aligned with public value. **Strong legal, regulatory, and institutional foundations are essential to ensure AI systems remain safe, equitable, and accountable throughout their lifecycle.** Many countries lack health-specific data laws and rely on general data protection frameworks that do not reflect the elevated sensitivity, risk, and public health implications of AI-driven data use.

> “This isn't a technology challenge; it's a governance and capacity challenge.”

Gaps in health-specific legislation

General data protection laws often prioritize privacy without providing mechanisms for responsible public-interest data use. This creates a “privacy chill” where data sharing stalls, pushing developers toward low-quality or unrepresentative datasets, and impeding the development of equitable AI solutions.

When proportionate and predictable, regulation accelerates AI adoption by clarifying expectations for safety, transparency, and accountability. Effective regulatory frameworks reflect public values, establish safeguards, and enable trusted innovation. Misaligned or inconsistent regulations across jurisdictions create compliance burdens and slow system-wide adoption.

Continuous evaluation, monitoring, and stress-testing

AI systems are dynamic and require oversight mechanisms that extend beyond one-time approvals. Evaluation frameworks must incorporate ongoing validation, post-market surveillance, risk-tiering, and defect detection to ensure performance remains stable across populations, environments and time.

Institutional capacity and technical expertise

Governments need regulatory agencies capable of **auditing AI systems, reviewing model documentation, analyzing lineage, and detecting bias.**

> "Good laws exist on paper, but without regulatory capacity they remain unimplemented."

Capacity gaps extend across ministries, regulators, health facilities, and frontline workers. Without investment in legal, statistical, and technical expertise, even strong laws remain unenforced.

3. Strategic alignment and purpose-driven AI

Purpose clarity is essential for aligning AI with public health outcomes. Without a **defined objective,** whether to strengthen the health system or improve population health, AI adoption risks drifting toward tool-centric experimentation rather than structural impact. Purpose-driven design emphasizes long-term population outcomes through centering the focus back on health goals, governance mechanisms, and systems-thinking approach rather than isolated technological deployments.

Aligning purpose and incentives for AI-driven preventive health

AI must be anchored in clearly **articulated public health goals.** Prioritizing technological capability over purpose leads to fragmented efforts that do not advance prevention, improve access, or address systemic gaps. Purpose acts as the organizing principle for governance, evaluation, data design, and diplomatic coordination.

Countries require a strategic horizon that recognizes the distinct timelines of service delivery, secondary prevention, and primary prevention. AI's most meaningful contributions will materialize when interventions are aligned with long-term visions, particularly in shifting systems from reactive care to prevention, where incentives remain structurally weak.

Preventive health generates limited immediate revenue compared to clinical treatment, constraining investment, and innovation in this domain. AI creates opportunities for real-time detection and personalized preventive strategies, but uptake depends on new incentive structures and governance models that recognize prevention as a public good.

4. Diplomacy, cooperation, and global alignment

AI's expansion across borders requires **governance approaches that move beyond national systems and toward shared standards, interoperable evaluation frameworks, and coordinated regulatory processes.** Because AI models, datasets, and vendors operate transnationally, fragmented regulatory environments increase risk, slow responsible adoption and disproportionate burden on low-capacity systems. Diplomatic alignment is emerging as a structural requirement for safe and equitable AI deployment at scale.

Cross-border convergence on standards and evaluation

Differences in approval processes, risk-tiering, model evaluation, and post-market monitoring create uncertainty for developers and health systems. **Coordinated mechanisms** allow technologies to be assessed once and recognized across jurisdictions, reducing duplication and accelerating safe entry into diverse markets.

Regional entities can play a critical role in **shaping coherent governance pathways.** Continental or regional frameworks such as those developed for health data governance support countries with varying capacities to adopt common principles, reduce fragmentation, and strengthen bargaining power with global actors.

Shared challenges across regions, including workforce shortages, multilingual delivery needs, and variable data quality, open opportunities for joint problem-solving and knowledge exchange. Collaboration enables countries to build capacity, design fit-for-purpose solutions, and develop governance models that reflect diverse realities.

Multilateral mechanisms for AI governance

Multistakeholder platforms and global networks provide convening power, shared evaluation tools, and common safeguards. These mechanisms help standardize expectations around safety, transparency, and oversight, supporting states with limited regulatory capacity to engage with rapidly evolving technologies.

5. Trust, institutions, and public engagement

Trustworthy institutions form the foundation for effective AI systems in health. Public willingness to share data, clinicians' acceptance of decision-support tools, and governments' ability to implement AI on a scale depend on transparent governance, accountability structures, and mechanisms that reflect societal expectations. Without trust, even technically capable systems face resistance, underuse, or outright failure.

Institutional credibility and AI governance capacity

Trust is shaped by whether institutions demonstrate competence, fairness, and accountability. Weak or opaque governance erodes public confidence, limiting data availability, and undermining AI performance. Institutions must be equipped to govern AI across the full lifecycle, i.e. evaluation, monitoring, and redress.

Avoiding extractive data practices

Concerns arise when data flows primarily benefit external actors without improving local health outcomes. **Governance frameworks** must prevent extractive dynamics in which low-resource settings supply data without receiving commensurate value, capacity, or autonomy in return.

AI's potential to strengthen health systems depends on **addressing structural constraints** that predate the technology itself. Fragmented data infrastructures, inconsistent identifiers, limited validation capacity, and weak interoperability remain core barriers to safe and equitable deployment. Without a **rights-based governance approach,** transparent evaluation mechanisms and institutions capable of monitoring systems over time, AI risks amplifying existing inequities and introducing new forms of misclassification and exclusion. Effective adoption requires **investments in data quality, workforce capability, legal regulatory safeguards** that support scaling of AI solutions without causing unintended harm to the public, and system designs that reflect actual clinical and public health environments.

Public voice and participatory mechanisms

Patient assemblies, citizen juries, and community oversight mechanisms can strengthen legitimacy and align AI deployment with societal values. These structures ensure that public expectations around inclusion, privacy, benefit-sharing, and risk tolerance are integrated into governance frameworks. People share data when they see value returned to them or their communities. AI ecosystems must therefore embed reciprocity through improved services, better health outcomes or transparent benefit-sharing. Without this, data withholding and distrust are likely.

In summary: cooperation as a governance imperative

Scan here to watch the full webinar:

Diplomatic cooperation will be essential as AI models, datasets, and vendors operate across borders. Convergent regulatory pathways, shared evaluation frameworks, and regional governance mechanisms can reduce fragmentation and strengthen safety and accountability. Trustworthy institutions and public engagement must remain central, ensuring that data use aligns with societal expectations and that benefits are returned to the people whose data fuel these systems. Building AI-enabled health systems is not a technical exercise, but a governance challenge that demands clarity of purpose, long-term vision, and coherent collective action.

Key recommendations and insights for stakeholders

Stakeholders of the AI for Health ecosystem

The AI for health ecosystem represents a convergence of disciplines and roles, bringing together different partners to enable the efficient adoption of AI, spanning the full continuum from AI tool design to the evaluation of health outcomes. Understanding the role each stakeholder plays, as well as their interdependencies and linkages with one another, is therefore essential (Figure 1).

Figure 1. **Mapping of diverse stakeholders across the AI for Health ecosystem.**

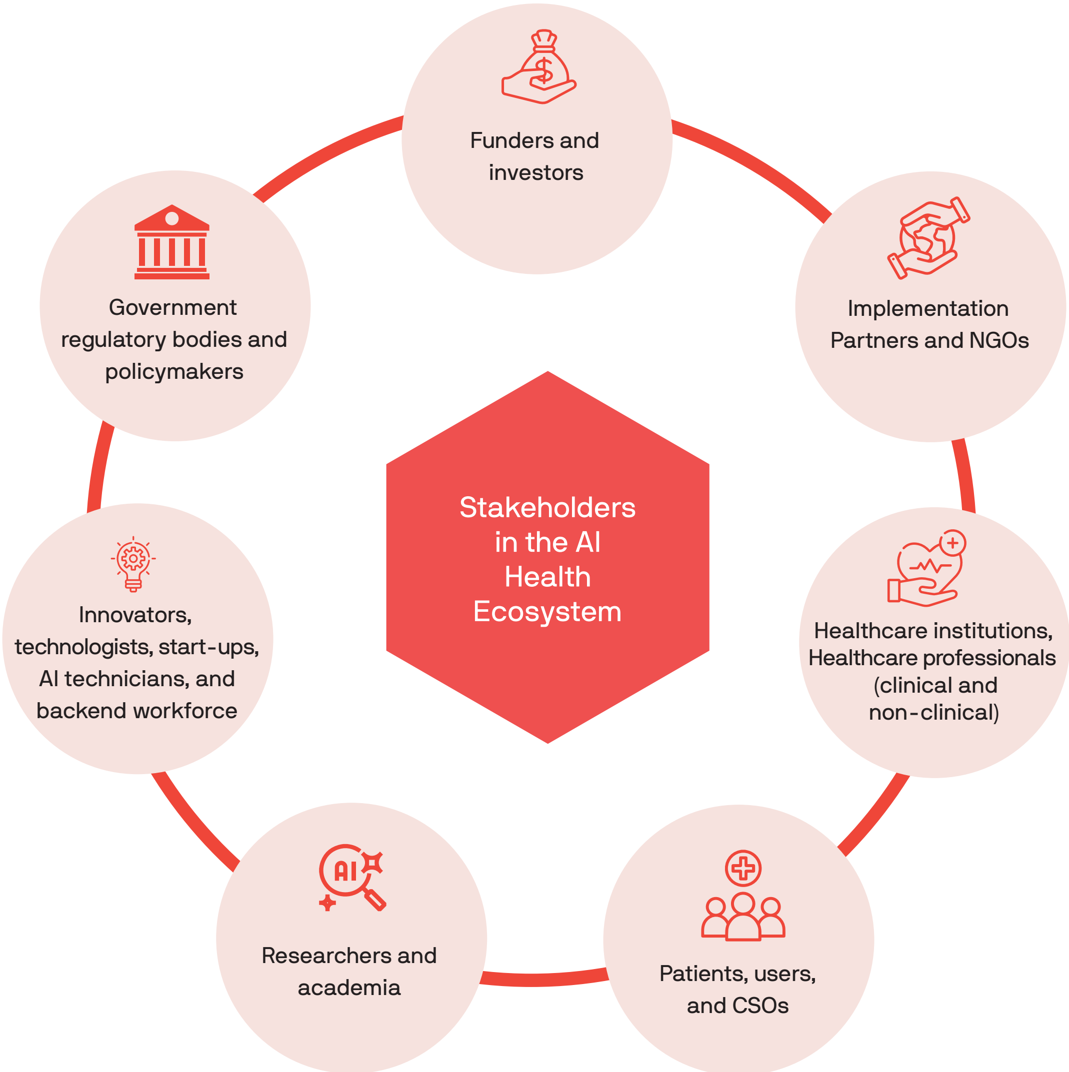

This mapping is broad by design, with the understanding that it may not represent the full universe of relevant actors. These stakeholders interact both linearly and non-linearly, directly or indirectly influencing one another across the ecosystem. For public health systems to fully leverage AI, the recommendations of this dialogue series must be interpreted and adapted relative to the position each stakeholder occupies within the ecosystem.

Stakeholder categories

Funders and investors: Financiers of both health systems as well as AI tools across different stages from development to pilots or large-scale roll out. These can be philanthropies, bilateral/ multilateral banks, venture capitalists, or CSRs.

Governments officials, policymakers and regulatory bodies: Together, these actors mediate the governance frameworks through which AI is applied within health care. Government officials set national health system priorities, shape health budgets, and drive the AI for Health agenda in their respective countries. Policymakers develop legislation and laws that support the health agenda, while regulatory bodies provide the checks and balances necessary to clarify boundaries and safeguard the system.

Innovators, technologists, start-ups, AI technicians, and backend workforce: This group encompasses developers and engineers, as well as the technical support staff responsible for data annotation and the maintenance of AI infrastructure.

AI researchers and academia: This category includes think tanks, independent researchers, universities, and academic circles that empirically assess the impact of AI in health, identify gaps in evidence and practice, and determine impact outcomes of AI projects.

Implementation bodies and nongovernmental organizations (NGOs): They facilitate coordination and alignment across the ecosystem, often bridging the gap between policy, technology, and on-the-ground delivery. This includes development partners and grassroot organizations.

Health care institutions and health care professionals (clinical and non-clinical): This group spans hospitals, pharmacies, diagnostic centers, and medical schools, as well as the full range of health workers including doctors, frontline workers, allied health professionals, support staff, and medical students.

Patients, users, and civil society organizations (CSOs): Communities and individuals who are the end users of AI-enabled health systems directly affected by the intersection of AI technology and health care service delivery.

Stakeholder-wise recommendations

A. Funders and investors

1. **Prioritize foundational systems for AI:** In-country governments must invest in the infrastructure that enables AI to function effectively. This includes building robust digital infrastructure, developing data platforms, subsidizing compute capacity, funding large-scale data collection initiatives, and strengthening workforce capabilities. Establishing these foundations is what will make AI adoption viable, scalable, and sustainable in the long run.
2. **Establish clarity on investment areas:** External funders including philanthropies, CSR funds, and venture capitalists should direct their investments toward the innovation space, focusing on creating an enabling environment and supporting the deployment of promising innovations. Critically, there needs to be clear delineation of who funds what, with active coordination mechanisms in place to ensure that external investments complement rather than duplicate existing government grants and public funding streams.
3. **Create predictable, milestone-based funding mechanisms:** Replace ad hoc grant cycles with rules-based processes that tie disbursements to milestones and measurable outcomes. This will allow health systems and innovators to plan for long-term adoption with confidence and stability.
4. **Calibrate funding instruments to the stage of development:** Use grants for pilots, blended finance models for infrastructure, and outcome-based contracts for interventions with proven population-level impact.
5. **Fund proposals with potential for sustainability:** Evaluate proposals not just on innovation, but on ecosystem fit, governance structures, and long-term sustainability. Prioritize tools aligned with national health strategies with built-in interoperability and avoid funding solutions that require continuous donor support.
6. **Mandate full cost modelling and total value assessment:** Require feasibility studies to go beyond license fees and account for data infrastructure, training, governance, and maintenance, alongside the cost of inaction at both micro and macro levels.
7. **Demand more from partners:** Hold private and philanthropic co-investors to a higher standard. Beyond capital, require them to contribute expertise, in-country training, and local capacity-building as conditions of partnership.
8. **Close the gap between innovation grants and procurement:** Extend funding horizons beyond innovation grants. Support mechanisms that link promising tools to long-term procurement pathways and ecosystem incentives, ensuring that successful pilots have a route to sustained system adoption.
9. **Condition funding on validation, data quality, and equity:** Require rigorous, transparent evaluation as a condition of funding. Invest in data infrastructure that enables performance assessment across diverse populations, finance data partnerships, and federated learning pilots to address the skewed datasets that drive inequitable AI outcomes.
10. **Keep equity non-negotiable:** Build equity safeguards directly into funding criteria. Require rural-first considerations, subsidies for underserved populations, and evidence of consultation with diverse end users before disbursing funds.
11. **Invest in change management and multi-stakeholder engagement:** Treat organizational change, training, and multi-stakeholder governance as a core funding line, not an afterthought. These are what determine whether a deployed tool is utilized effectively.
12. **Enable cross-border alignment and cooperation:** Invest in regional health data governance and AI evaluation frameworks. Support multilateral mechanisms that reduce regulatory fragmentation, enable shared safeguards, and facilitate benefit-sharing models that prevent extractive use of LMIC data.

B. Government officials, policymakers, regulatory bodies

1. **Mainstream AI into core health policy:** Move AI out of innovation silos and integrate it into mainstream health policy, procurement frameworks, and clinical leadership structures so it becomes part of how the system works, not a parallel experiment.
2. **Design scale pathways before pilots begin:** Define the path from proof-of-concept to embedded practice before a pilot begins. Establish government-led scaling hubs, financing mechanisms, and governance structures at the outset. The transition from pilot to system adoption must be pre-planned and not improvised.
3. **Build a regulatory environment that enables responsible innovation:** Develop rights-based, health-specific legal frameworks that protect sensitive data while enabling responsible use. Adopt risk-proportionate approvals, invest in post-market surveillance, and pursue mutual recognition with peer regulators. Build institutional capacity for continuous evaluation, stress-testing, and algorithmic auditing, and invest in specialist training and toolkits that allow oversight bodies to keep pace with deployment.
4. **Create structured pathways for safe experimentation:** Complement pre-market requirements with sandbox environments, phased rollouts, and clearly defined performance thresholds. Structured experimentation allows governments to test innovation responsibly without exposing populations to premature risk and give developers the predictability they require.
5. **Legislate clear data protection, ownership, and sovereignty frameworks:** Enact enforceable rules on patient privacy, data rights, and cross-border data flows developed in alignment with local implementation realities and through international cooperation. Communities without credible safeguards will resist AI, and systems without them face mounting legal and reputational risk.
6. **Invest in digital infrastructure and assess readiness:** Invest in the digital backbone that makes AI viable such as interoperable systems, stable patient identifiers, and consistent record-keeping. Conduct thorough readiness assessments before committing deployment: infrastructure gaps in connectivity, power, and interoperability must be addressed alongside digital literacy and workforce training.
7. **Require vendors to demonstrate affordability and switching feasibility:** Scrutinize the true cost of AI adoption for public systems. High switching costs and opaque pricing models deter adoption. Require vendors to demonstrate cost equivalence or savings relative to current systems and ensure business models are affordable at a public sector scale.
8. **Lead public communication and institutionalize accountability:** Communicate clearly about how and why AI is being used in public health systems. Resource oversight institutions such as data commissioners, ethics boards, and regulatory agencies, and formalize community voice through citizen assemblies, patient forums, and transparent mechanisms for consent, redress, and accountability.
9. **Align with regional and multilateral AI governance frameworks:** Engage actively in regional and multilateral governance frameworks. Align national AI strategies with international mechanisms that offer shared standards, evaluation tools, and technical support. Negotiate reciprocal data arrangements that ensure health data from LMICs is not misused and extracted without commensurate benefit.
10. **Recognize and protect the “invisible” AI workforce:** Extend workforce policy to cover the full AI labor chain, including the data labelers, annotators, and backend workers whose contributions are frequently invisible. Mandate fair compensation and include this workforce into national AI workforce strategies.

C. Innovators and technologists, AI technicians, and backend workforce

1. **Define the problem before building the solution:** Ground every product decision in a clearly identified gap in health service delivery and resist retrofitting existing technology onto workflows it was not designed for.
2. **Apply human-centered design principles to co-design with frontline workers and build for local realities:** AI solutions must be developed collaboratively with end users, such as frontline workers, clinicians, and nurses, to ensure the tool can meet their needs and be aligned with the cultural and linguistic contexts for sustainable adoption.
3. **Build for real-world infrastructure, not ideal conditions:** Design for real-word settings, i.e. limited connectivity, shared devices, and erratic power supply as these are the norm in most deployment settings. Solutions that do not account for these factors will stall and subsequently face challenges during implementation.
4. **Integrate into health systems, not alongside them:** Build tools that connect to existing systems and real clinical workflows and not standalone applications. Validate tools in the most demanding environments first and prioritize intuitive, user-friendly interfaces, offline functionality, and minimal training requirements.
5. **Choose technology modalities strategically:** Match technical approach to the task. Traditional ML remains more reliable for narrow, high-accuracy tasks in low-resource settings. Generative AI can aid in decision support but carries regulatory, performance, and cost challenges that must be evaluated before deployment.
6. **Be transparent about what your tool does and monitor it continuously:** Document intended use, known limitations, risk profiles, and the evidence base clearly. Build real-time performance tracking, post-deployment audits, and model recalibration mechanisms into the product from day one. AI systems that perform well at launch can drift, and catching issues early is what separates responsible deployment from harm.
7. **Ensure fair recognition and compensation for the full AI workforce:** Data labelers, annotators, and backend workers are foundational to AI performance. Advocate for their fair pay, visibility, and inclusion in workforce protection frameworks, and hold your own organizations accountable for equitable practices.

D. AI researchers and academia

1. **Prioritize validation frameworks and locally representative datasets:** High-quality, localized, and well-structured data is the foundation of trustworthy AI. Build frameworks that enable rigorous performance assessment across diverse populations, treat data infrastructure as a scientific output, and generate locally grounded datasets-especially for historically underrepresented communities.
2. **Make bias identification and mitigation a research discipline:** Systematically quantify where AI models reproduce inequities in diagnosis and clinical support and develop methodologies that make bias mitigation standard practice.
3. **Set and uphold transparent evidence standards:** Define what AI tools claim to do and under what conditions, challenge overstated performance claims, and use peer review as an active safeguard against misleading claims entering clinical environments.
4. **Conduct evaluation studies of deployed AI tools to measure actual impact:** Support evidence-based decision making by conducting studies validating and assessing if the AI tools genuinely improved care delivery, health outcomes, or system efficiency.

5. **Extending research into deployment and post-market performance:** Develop auditing methodologies, drift detection tools, and recalibration frameworks, and embed equity impact metrics into study design to assess how AI affects access and quality for underserved populations.
6. **Break disciplinary silos and inform context-sensitive governance:** The most consequential AI research in health happens at the intersection of clinical expertise, technical innovation, ethics, and community knowledge. Collaborate with clinicians, policymakers, and community representatives to produce findings that inform adaptive, context-sensitive governance and not just universal frameworks. Disseminate findings in formats that policymakers from diverse settings can act on.

E. Implementation bodies and NGOs

1. **Translate policy into practice at the last mile:** Bridge the gap between national AI strategies and frontline delivery by using implementation expertise to adapt tools, train workers, and troubleshoot deployment in the low-resource, high-complexity settings where most health AI fails.
2. **Anchor implementation in community trust and local knowledge:** Leverage existing community relationships to identify contextual barriers, build acceptance, and ensure that AI tools reflect the languages, norms, and care pathways of the populations they are meant to serve.
3. **Document learnings, implementation realities, and feed them upstream:** Systematically capture what works and what fails in deployment, and make that evidence available to developers, funders, and policymakers so that future tools are designed to withstand real-world constraints.
4. **Protect frontline workers from AI-related burden and displacement:** Monitor the impact of AI tools on health worker roles, workloads, and confidence and flag deployments that increase burden, create confusion, or undermine the professional standing of CHWs and allied staff.
5. **Build local ownership and reduce dependency on external implementation support:** Design implementation programs that transfer skills, tools, and decision-making authority to local institutions and health workers as sustainable AI adoption hinges on local capacity, and not permanent external facilitation.

F. Healthcare institutions and healthcare professionals (clinical and non-clinical)

1. **Participate in design, not just deployment:** Bring frontline and clinical knowledge into the co-design processes from the outset. Clinicians and allied health workers who engage early in design can influence whether a tool reduces burden rather than adding to it, so that tool aligns with their workstreams and doesn't disrupt it.
2. **Clinician's judgement remains at the center of care:** Use AI as a decision-support tool that enhances professional expertise and relationships with patients. Maintain the right and responsibility to question, override, and escalate beyond any AI output where clinical assessment demands.
3. **Demand workflow fit, consistency, and long-term commitment:** Push back on fragmentation and frequent system changes. Health care workers will only invest in AI tools if they trust those tools are stable and won't be replaced in the short term. Insist on continuous testing against real clinical environments and hold vendors accountable for tools that disrupt clinical workflow.
4. **Integrate AI literacy into professional development for the health care workforce:** Institutions should integrate AI training into onboarding, continuing education, and clinical curricula, providing mentorship

that builds confidence rather than surface-level familiarity. Where possible, ensure frontline health care workers have access to the devices they instead of relying on personal devices.

5. **Monitor performance, maintain feedback loops, and lead multi-disciplinary review:** Establish systems for tracking AI performance after deployment and capture and act on frontline experience and feedback. Real-time monitoring, regular audits, and clear patient safety reporting protocols allow institutions to catch drift, bias, or declining accuracy before it affects care.

G. Patients, users, and CSOs

1. **Understand how AI is being used to deliver care:** AI supports rather than replaces the clinician in health care delivery. Patients and users need to understand how AI is being used in their treatment processes. The “human in the loop” is not a non-negotiable feature; rather commitment to safety and dignity.
2. **Know, assert, and protect data rights:** Patients and users are the owners of their own health data. They must be made aware of how their data is collected, used, and shared, including training AI models, and withhold consent where meaningful protections are absent.
3. **Expect tools designed for real-world situations and report when they are not:** AI in health should improve access and quality for everyone, including those in rural, low-income, or marginalized communities. If a tool does not reflect the local realities of its users, it is a design failure. Communities are entitled to report when tools do not work for them, and to expect that feedback leads to change.
4. **Engage in AI governance and hold institutions accountable:** Citizens can participate in assemblies, patient forums, and community consultations to demand transparent, accessible mechanisms for raising concerns and seeking redress when AI causes harm.

Closing note

This compendium is the culmination of the **AI and Public Health: Expert Dialogues on Strategy, Governance, and Ecosystem Transformation series,** which has been a deliberate and sustained, multi-stakeholder effort to surface what is working in AI for health, what remains unresolved, and what the path forward requires. Cross-cutting across five thematic areas—scaling AI for health impact, financing, workforce, regulation and ethics, and global diplomacy—the series brought together governments, implementers, researchers, clinicians, regulators, community representatives to share evidence, learnings, debate trade-offs, and develop shared understanding.

What we have is not a singular perspective, but a 360-degree view of the AI for health landscape, right from the ground-up to the decision-making level. This extensiveness is intentional. The challenges AI poses to public health systems are systemic, and so must be the responses.

The priorities, recommendations, and actionable pathways compiled here are designed to serve diverse audiences. For countries at early stages of AI adoption, they offer practical guidance on where to start, what to anticipate, and how to build for scale rather than for pilots alone. For policymakers and technical bodies shaping regional or global frameworks, they surface the cross-cutting governance questions that single jurisdictions cannot resolve alone. For practitioners and researchers, they provide a grounded, evidence-informed reference point for a field that is moving faster than any one institution can track.

Taken together, these chapters do not offer a blueprint since the diversity of health systems, resource contexts, and governance traditions makes that neither possible nor desirable. What they offer instead is a foundation rooted in shared principles and lessons, as well as a clear account of where the work of making AI safe, equitable, and effective for public health still needs to go.

PATH

15th Floor, Dr. Gopal Das Bhawan
28, Barakhamba Road
New Delhi - 110001