مؤسسة قطر
Qatar Foundation

WISH
مؤتمر القمة العالمي للابتكار في الرعاية الصحية
World Innovation Summit for Health

8
النسخة
EDITION

Policy Report

NOVEMBER 2026

Artificial intelligence across health systems

A report on the safe deployment, scaling and governance of artificial intelligence in health systems

Amish Acharya
Annas Gamaty
Peter Howitt
Deidre Thompson
Ahmed Elmagarmid
Ara Darzi

Suggested reference for this report: Acharya A, Gamaty A, Howitt P, Thompson D, Elmagarmid A, Darzi A. Artificial intelligence across health systems: A report on the safe deployment, scaling and governance of artificial intelligence in health systems. Doha, Qatar: World Innovation Summit for Health, 2026.

ISBN: 978-1-913991-45-6

CONTENTS

2 Foreword

3 Executive summary

5 Section 1. Introduction

7 Section 2. The expanding human–AI interface in health

11 Section 3. Current policy and governance

14 Section 4. Establishing trust

17 Section 5. Deployment

21 Section 6. Governance

23 Section 7. Evaluation

26 Section 8. Putting EDGE AI into practice in Qatar

30 Section 9. Recommendations

34 Appendix. Key AI source documents

39 Abbreviations

41 Acknowledgments

42 References

FOREWORD

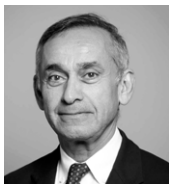
Artificial intelligence (AI) is no longer an emerging prospect in health but an operational reality. More than 1,200 AI-enabled medical devices have now been authorized by the United States Food and Drug Administration, and comparable products are reaching markets and clinics across many regions. From radiology to general-practice consultations, AI tools can improve patient-level outcomes and experience, while relieving pressure on overstretched health services. However, as we develop tools, including generative and multiagent systems directly accessible to patients, there is a need to ensure safe, equitable and durable deployment which, in many contexts, is not yet in place.

This report is written for those responsible for that task – ministers and ministries, heads of service, clinical and informatics leaders, regulators, funders and the non-governmental organizations that support them. All policymakers are having to react to an evolving environment. The speed of discovery is vastly outpacing our capacity to govern health AI systems that are both safe and inclusive.

Authorization is not assurance. Meeting technical requirements is only part of the life cycle of tools capable of continuous and dynamic iteration. Safe deployment and scale depend on downstream factors and rely more on the receiving environment than on the incoming innovation.

Innovation can be transformative rather than evolutionary. In words often attributed to Henry Ford: *"If I had asked people what they wanted, they would have said faster horses."* However, such transformative innovation also brings with it new risks. Five years after Ford's Model T first came off the production line, approximately 33 people died for every 10,000 vehicles on the road – more than 20 times the most recent US rate. The key is to not hinder innovation but rather put in place the necessary frameworks and policies that enable it to thrive to the benefit of all people.

As we sit amidst another great technological revolution, we must convert principles into practice and ensure that the innovation which defines our generation transforms our health systems for good.





Professor Lord Ara Darzi
Director of the Institute of Global Health Innovation and Paul Hamlyn Chair of Surgery at Imperial College London; Executive Chair, WISH

Ahmed Elmagarmid
Executive Director of the Qatar Computing Research Institute, HBKU; Champion of AI HUB; Vice Chair of the Qatar Science and Innovation Foundation

EXECUTIVE SUMMARY

Artificial intelligence (AI) is moving rapidly from research into routine healthcare. AI systems are now used across the patient pathway, from clinical workflow tools to patient-facing applications and emerging agentic systems working autonomously on tasks. While technological development has accelerated, the capability of health systems to deploy, monitor and govern these technologies consistently has not kept pace.

The decisions made now will determine how AI is integrated and whether it will improve healthcare in the near future. This report examines the conditions required to translate AI from technical success into sustainable impact for health systems. The report concludes that AI is more likely to augment or supplement clinical workflows rather than replace them entirely. Also, the impact of AI in health will depend more on the strength of the health system, rather than on the sophistication or innovation of the AI algorithms.

Across jurisdictions, policy approaches are increasingly converging around common principles, including risk-proportionate regulation, transparency, human oversight, validation against intended use and post-market monitoring. However, important differences remain in how countries manage adaptive models, generate real-world evidence, commission AI technologies and define accountability for deployment. Significant gaps also persist, particularly for direct-to-consumer AI, organizational implementation and post-deployment oversight.

This report examines the current policy and governance landscape and identifies four interdependent components of successful AI adoption which have been compiled into a framework – EDGE AI:

- **Establishing trust:** The underpinning public and professional acceptance of the AI system is a prerequisite to scale. It involves moving from individual accountabilities to system-level structures that can foster trust, such as data privacy and cybersecurity measures.
- **Deployment:** Determines whether AI can be integrated into clinical pathways. Its dependencies range from data foundations of the model to workforce strategies, procurement mechanisms and culture.
- **Governance:** Establishes lines of accountability and regulatory mechanisms required throughout the AI life cycle. A risk-based assessment ensures that necessary oversight is balanced with meeting healthcare needs.
- **Evaluation:** Provides evidence of real-world performance, value and safety. Defining critical performance metrics and ensuring real-world evaluation supports recalibration, retraining or withdrawal where performance deteriorates.

Qatar offers a useful test case because the country has invested in foundations aligned with the EDGE AI framework. The country's experience shows that the framework is successful in deploying, governing and continuously evaluating AI

while maintaining trust, safety and equity, and that AI readiness is built progressively through investment in health-system capabilities suited to national context and priorities.

This report recognizes the varying capabilities of different health systems. Rather than prescribing a single model, the recommendations consider a country's level of AI maturity and propose eight practical actions for countries at different stages of AI adoption. This allows governance, evaluation and deployment capabilities to develop progressively while maintaining patient safety and public confidence.

The recommendations cover:

- Recommendation 1. Equity by design
- Recommendation 2. Deliberate human–AI interaction and trust
- Recommendation 3. Govern the deployment, not just the product
- Recommendation 4. Procurement for outcomes, interoperability and adaptability
- Recommendation 5. Patient-facing tools governed by function and risk
- Recommendation 6. Named accountability
- Recommendation 7. Resource monitoring with thresholds and set deadlines
- Recommendation 8. Independent evaluation capability

The policy recommendations focus on strengthening public trust, enabling responsible deployment, embedding real-world evaluation into routine practice, establishing proportionate governance, and supporting international collaboration. Collectively, these measures aim to provide policymakers with a practical framework for deploying AI safely, generating evidence of value, and enabling sustainable adoption across diverse healthcare systems.

SECTION 1. INTRODUCTION

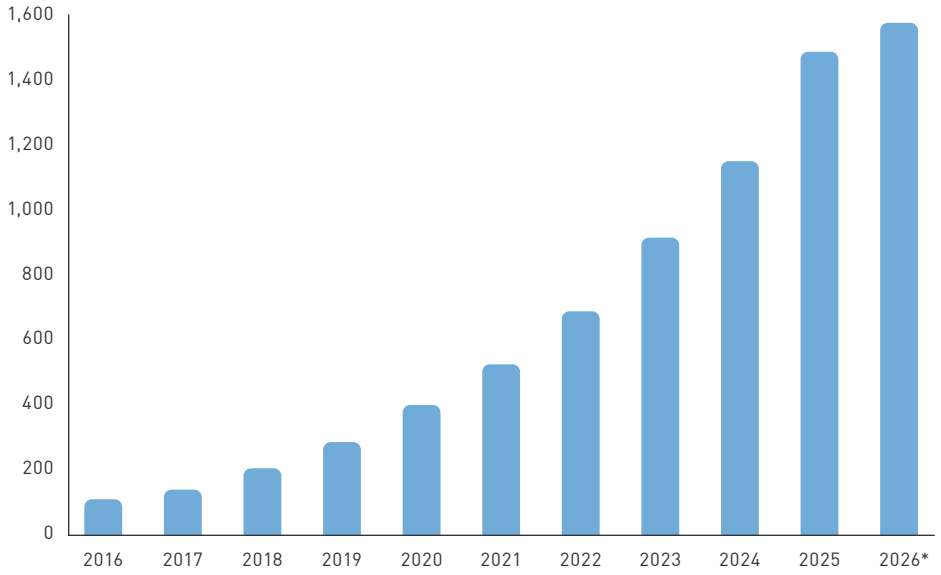
1.1 CONTEXT AND RATIONALE

Longer lives do not always equate to healthier ones. Health systems worldwide are facing rising demand, which is outstripping workforce capacity and budgets.

Traditionally, healthcare has been designed around physical infrastructure – hospitals, clinics and practices – areas that define patient flow, access to health information, and outcomes. In many low- and middle-income countries (LMICs) these limitations are particularly noticeable, as distance to the nearest health facility with appropriate capacity can be prohibitive.

Against this backdrop, the number of authorized AI-enabled medical technologies has increased substantially, with record numbers reported year-on-year, as shown in Figure 1.

Figure 1. Cumulative artificial intelligence/machine learning devices authorized by the US Food and Drug Administration each year



*Up to 30 March 2026

Source: United States Food and Drug Administration (FDA) (2026a).[1]

Many AI tools offer novel insights and understanding from existing and widely available tests. For example, AI applied to handheld electrocardiogram (ECG) recordings can identify patients likely to have abnormalities that would otherwise require echocardiography to detect.[2] Such approaches can help to decentralize care, and enable earlier diagnosis, triage and intervention while supporting a shift towards prevention. However, AI tools are reaching patients and evolving faster than the systems designed to evaluate and oversee them.

The emergence of foundation models (AI systems trained on large data sets) with multiple downstream applications that can be rapidly adapted and adopted is challenging traditional regulatory approaches.[3] Conversational AI is increasingly becoming a first source of health information. An analysis of 500,000 de-identified Microsoft Copilot health conversations found that while most queries were informational, nearly one in five involved users describing personal symptoms, test results or conditions without formal clinical oversight.[4]

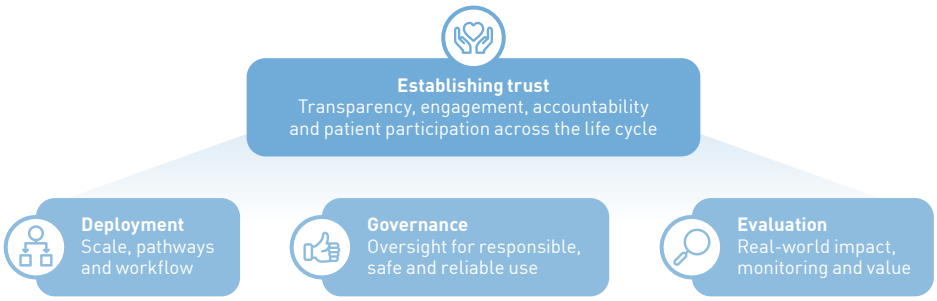
As a result, a gap is emerging between regulatory approval and real-world performance. For example, despite being used across hundreds of hospitals, a widely used sepsis early warning model identified only about one-third of sepsis cases when externally validated at a large academic health system.[5]

This challenge is amplified for large language models (LLMs) – AI systems trained to understand and generate human-like text – where performance depends on how users interact with the model. One study found that diagnostic accuracy was up to 63.6 percent lower when prompted by members of the public as opposed to researchers in simulated settings,[6] reinforcing the need for stronger governance and ongoing evaluation.

It is increasingly evident that health systems realize the greatest benefit from AI when governance is trusted and AI is integrated into clinical and operational workflows, supported by continuous real-world evaluation. While approaches differ between countries, there remains a need for practical frameworks that support safe and effective implementation across health systems with varying levels of maturity and resources.[7]

The EDGE AI framework (see Figure 2) recognizes that successful AI depends on the readiness of health systems, rather than simply the availability of AI technologies.

Figure 2. EDGE AI: A framework to guide AI adoption



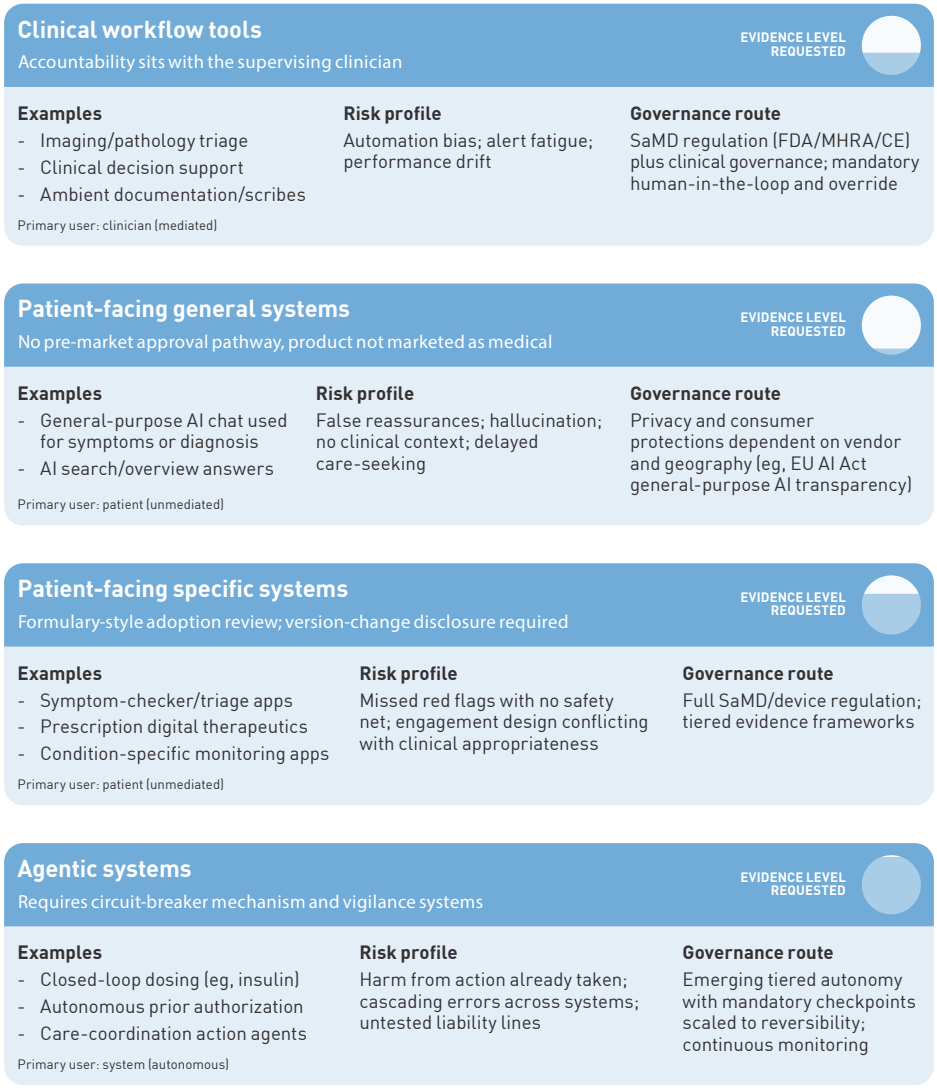
The EDGE AI framework provides policymakers with a practical approach for supporting the safe and effective adoption of AI. The framework encompasses digital infrastructure, interoperable health information systems, data quality and governance, workforce capability, institutional leadership, financing mechanisms, regulatory capacity and continuous evaluation. Together, these capabilities determine whether AI can be deployed safely, equitably and sustainably.

SECTION 2. THE EXPANDING HUMAN–AI INTERFACE IN HEALTH

The human–AI interface in healthcare is rapidly evolving. AI is changing how patients and doctors interact with each other, how they view, process and act upon medical data, including test results.

In the short term, AI is more likely to extend, augment or supplement clinical workflows rather than replace them entirely. The maturity of AI applications differs widely across use cases, as do the evidence thresholds, risk profiles and the governance needed to translate these applications from pure research to patient impact. Figure 3 provides an overview of the current AI systems in healthcare.

Figure 3. A taxonomy of healthcare AI uses



Sources: Adapted from European Union (EU) (2024a & 2024b);[8,9] Medicines and Healthcare products Regulatory Agency (MHRA) (2024a);[10] American Medical Association (2026).[11]

While clinical workflow tools such as imaging and diagnostic support are comparatively mature, triaging tools are new but growing in prevalence. Agentic systems have the least settled evidence and governance but are of increasing interest. The concept of 'agent hospitals' is emerging in countries such as China, where it facilitates real-world integration of AI tools by simulating entire patient journeys.[12]

AI across the patient pathway

As the use of AI becomes more widespread, it is becoming more important to appreciate user interactions to avoid unintended consequences, including unregulated use. For example, a recent survey of more than 2,000 healthcare professionals across 10 countries showed that, in the absence of organizational access, 64 percent of clinicians resort to using personal AI tools.[13] Healthcare systems need to ensure that regulation facilitates appropriate use.

2.1 CLINICIAN-WORKFLOW TOOLS

The World Health Organization (WHO) estimates that, by 2030, there will be a deficit of more than 10 million health workers.[14] AI can help fill this gap, freeing clinicians from labor-intensive administrative tasks or supporting decision-making by aiding clinical processing, comprehension and presentation of multimodal health data. For instance, outside of the clinical encounter, AI systems such as Docman AI or Epic Penny can read complex clinical correspondence, extract relevant information, and either route it appropriately or generate a structured response.[15,16] However, systems that support decision-making – requiring reasoning over complex clinical data to generate insights, recommendations or alerts – pose a greater risk (see Box 1).

BOX 1. AI CO-CLINICIAN

In April 2026 Google DeepMind launched a new initiative, AI Co-Clinician. Moving on from medical AI models such as Med-PaLM that were trained to answer medical examination questions, AI Co-Clinician explores the capabilities of AI agents to augment the clinical team, including interacting with patients under human supervision.[17]

AI Co-Clinician is currently being tested in healthcare systems in the US, India, Australia, New Zealand, Singapore and the United Arab Emirates in two task groups: clinician assistance, and patient navigation. In the former, when presented with curated primary care responses, AI Co-Clinician provided fewer critical errors, and its responses were subjectively preferred by clinicians compared to two AI systems widely used by physicians. It also showed higher accuracy when responding to simulated medication or management queries. Leveraging the capabilities of Gemini and multimodal processing underpinning Project Astra, AI Co-Clinician can use live audio and video to engage with patients within telemedical settings. Tested in simulated patient scenarios, the system was able to correct an asthma inhaler technique or guide a patient through physiotherapy techniques.

However, the system was less adept than physicians at identifying serious or life-threatening 'red flag' symptoms within consultations. This highlights the need to: maintain human oversight in clinical decision-making with current medical AI technology; understand accountability with decision-making; and ensure that AI models' operational safeguards are rigorously maintained.

Source: Shah et al. 2026.[18]

In these scenarios, failure to fully understand and design the interaction between clinicians and the AI tool can lead to significantly negative outcomes, usually including two interrelated biases: automation bias and alert fatigue. The first represents a tendency to overly rely on AI outputs, even when contrary evidence is present; the second occurs when low-level alerts or continuous alarms lead clinicians to ignore the systems. Robust oversight of clinician–AI interactions – what is being surfaced, when, and to whom – is key to preventing these biases from manifesting unexpectedly following the use of AI.

2.2 PATIENT-FACING GENERAL AND SPECIFIC SYSTEMS

The fastest-moving and least-governed frontier is AI that patients use directly – including symptom-checkers, general-purpose conversational assistants used for health questions, wellness and companion applications, and emerging generative therapy chatbots. These tools can expand access and support self-management, particularly where professional capacity is scarce or inaccessible.[19,20,21] But they raise distinctive risks. Generalized models can have insufficient information about an individual's health status or use unverifiable sources when answering queries, and can lack the necessary safeguards.[22]

Much of this activity falls outside medical-device regulation. In the US, the Food and Drug Administration (FDA) has historically exercised enforcement discretion over symptom-checkers.[23] The first clearance of a device incorporating a patient-facing LLM was in December 2025 – even as the agency's Digital Health Advisory Committee began to examine generative mental-health and therapy tools.[24]

Oversight is emerging, albeit unevenly:

- Consumer protection inquiries into companion chatbots.
- State-level laws requiring disclosure that a user is interacting with AI, prohibiting chatbots from presenting themselves as licensed clinicians.
- Mandated safeguards for minors and crisis detection.[25,26]
- Patient-facing systems can also fall under the general wellness low-risk exemptions, unless the AI is making specific medical claims.[27]

Deployers and policymakers must adhere to key principles. Tools that make individualized health claims, and are accessible to the public, should be treated according to function and risk, not marketing category. Users should always know when they are interacting with AI, what its limits are, and whether other specialized

health tools are available. Safeguards must be in place for vulnerable groups, including children, people in crisis or those with lower literacy, especially when queries pertain to wellness and health product use (see Box 2).

BOX 2. COPILOT HEALTH

Based on the 50 million consumer health questions posed daily to Microsoft's products, in March 2026, Microsoft launched Copilot Health, a separate space within Copilot that offers direct-to-consumer health information. The system responds specifically to health-related queries. It provides clearly cited source material from sources considered to be credible, leveraging a partnership with Harvard Health.[28] The system also connects to real-time US provider directories to help patients navigate provider locations, available languages and insurance accepted.

Copilot Health also connects with repositories where users can upload data from wearable devices supplied by various vendors, and the health records of more than 50,000 US hospitals and provider organizations. Responses can be personalized or at least targeted to an individual's medical status.

Copilot Health has achieved ISO/IEC 42001 certification, and is currently only available in the US. There are detailed safeguards surrounding data privacy, and the ability to turn off personalization. However, data on the system's real-world performance are not yet available, including the consequences for patient care when sub-optimal answers are provided to users' queries.

Source: Gross et al. (2026).[29]

2.3 AGENTIC AI

Agentic AI refers to systems that can adapt autonomously and execute tasks with minimal human intervention to deliver defined clinical goals.[30] These types of AI exhibit the highest levels of decision-making, adaptability and self-sufficiency.[31] In many cases, these systems use multiagent frameworks to break down complex clinical tasks or to offer adversarial positioning within decision-making. For instance, teams have developed proof-of-concept agentic systems for automating radiotherapy plan optimization. Here a dosimetrist; an agent to manage the planning system; a physicist agent to evaluate dose; and a human agent enabling human input were together able to develop acceptable treatment plans for lung cancer patients, reducing mean dose and increasing target coverage compared to existing software.[32]

While fully autonomous clinical agentic systems are not widespread, agentic AI has moved beyond proof-of-concept in administrative, operational and research workflows, including patient scheduling, care co-ordination, revenue cycle management and evidence synthesis.[33] As use of this AI continues to spread, there is an increasing need to have clear lines of oversight, accountability and monitoring.

SECTION 3. CURRENT POLICY AND GOVERNANCE

In the last five years, health AI policy has accelerated, most notably between 2024 and 2025,[34] with a mixture of hardline regulation, supplemented by softer guidance from a growing body of issuers – including national and intergovernmental organizations – looking beyond development to deployment of AI health tools.

3.1 PRE-MARKET VALIDATION

Many countries use similar guidance, often supporting Good Machine Learning Practice and the value of intended purpose.[35] Many share a position on the need for high-quality and representative data sets, ground truth or reference standards and traceability. The Organisation for Economic Co-operation and Development (OECD) and WHO offer frameworks for the non-binding principles and validation considerations that most national regimes draw on, such as the need for data representativeness, bias mitigation, and human oversight.[36–38] The trend is for AI life cycles, not just specific performance, moving toward risk-based assessments of effectiveness and generalizability of models.

Markets such as the European Union (EU) and China use binding conformity assessments requiring third-party or registration review before market entry.[39–43] The US largely layers guidance onto existing pre-market pathways, and life cycle recommendations by the FDA remain in draft. In many African countries, specific AI regulatory frameworks are still in development. In early 2026, the Pharmacy and Poisons Board (PPB) published the *Guideline on Regulation of Medical Device Software in Kenya,* which incorporates AI tools – including the need for pre-market reporting on the source, data type collected, data pre-processing steps and parameters being extracted for continuous learning models.[44]

3.2 POST-MARKET SURVEILLANCE (PMS) AND MONITORING

The International Medical Device Regulators Forum Software as a Medical Device (SaMD) framework[45] has fostered alignment in PMS. AI-specific risks such as entrenching bias and data drift are important for developers, regulators and adopters to consider. These risks raise the need for co-ordinated monitoring mechanisms, with human oversight as a safety net to avoid the accumulation of errors. The following systems all use predefined metrics to facilitate early detection of degradation:

- WHO's continuous-monitoring language[36]
- Kenya's PPB post-market commitment[44]
- Medicines and Healthcare products Regulatory Agency's (MHRA's) adaptive-AI work[46,47]
- FDA's monitoring-plan expectation[1,48]
- Singapore's Artificial Intelligence in Healthcare Guidelines (AIHGle)[49]

The UK requires a robust PMS system for medical devices, including regular reporting (the frequency determined by risk), vigilance processes, and the collection of real-world data to monitor performance including safety.

The OECD's AI incident reporting framework – *Towards a Common Reporting Framework for AI Incidents* – and networked incident sharing also highlight the concern over AI-specific failures such as bias, drift or hallucination, which need prompt action if discovered.[50]

However, where future changes are required – for example, to update a continuously learning AI model – there are sharper divergences between geographies. The US Predetermined Change Control Plan (PCCP) enables pre-authorization for defined future changes without the need for resubmission to regulators, a process similar to the UK's MHRA and Health Canada.[48,51] The EU and China lean more toward a reassessment process, or in some cases, reregistration if changes are deemed significant.[40,41,52] Across all jurisdictions, incident reporting taxonomies for AI-specific harms are immature and non-standardized.

3.3 DISSEMINATION AND COMMISSIONING

Regulation of AI health tools is being increasingly positioned as the entry point for dissemination and scale. Many regions, including the US and EU, set out frameworks for deployment, relying on individual states to determine how this is conducted in practice.[53,54] The UK offers the AI and Digital Regulations Service (AIDRS) with a single-payer system that supports innovators and the healthcare system to bring safe, impactful AI and digital technologies onto the market and into routine use.[55] Through its living AIHGle framework, Singapore disseminates practical guidance to support the safe development and use of AI in healthcare, complementing the Health Sciences Authority's regulatory frameworks.[49, 56] There is growing consensus among territories that scaling AI depends on enabling infrastructure – interoperable data, update mechanisms, harmonized market access, and PMS – rather than on strict regulation.

3.4 SYNTHESIS

To maximize health AI value, policy approaches commonly address core principles such as transparency, risk proportionality, human oversight, validation against intended use, bias mitigation, and PMS. However, there are significant differences in how jurisdictions operationalize and enforce these principles. For example, approaches vary in how risk is defined across geographies, how health systems manage adaptive or continuously learning models (eg, pre-authorized updates or re-assessments), and whether formal mechanisms exist to link evidence generation with commissioning and reimbursement decisions – where system developers and healthcare providers are financially compensated for using AI models. Requirements for evidence submission, oversight and scrutiny also differ considerably.

Important gaps remain, and some may create regulatory blind spots as AI tools become more widely deployed in health systems:

1. Most regulatory frameworks focus on products before market entry, with limited attention to oversight of system-level use, organizational deployment, validation in practice and integration into care pathways.

2. Direct-to-consumer and open-source AI models that provide health-related information often fall outside the scope of existing regulatory frameworks.

3. Considerations specific to AI-enabled medical devices are addressed in some jurisdictions, but broader guidance is often lacking on implementation, governance and health system use.

Establishing scalable AI use in healthcare requires co-ordinated action across regulators, health systems, clinicians, developers and funders. Although implementation pathways vary by jurisdiction, successful adoption depends on clear regulatory processes, proportionate oversight and supporting infrastructure. Therefore, policymakers should focus on market authorization and on the conditions required for safe, effective and sustainable deployment at scale.

SECTION 4. ESTABLISHING TRUST

Trust in AI is a prerequisite for success. In a survey of surgeons from 98 countries, while perceived patient trust was not seen to be a barrier to AI adoption, 62.7 percent of respondents stated that clinical accountability and transparency were an increasing concern.[57] Paralleling the OECD's position, we consider trustworthiness as the foundation of the EDGE AI framework, which supports the safe and effective adoption of AI.

4.1 INDIVIDUAL ACCOUNTABILITY

Accountability remains a significant concern among clinicians relying on AI-generated insights in practice.[58] When an AI-informed decision causes harm, it is essential that the entity or person responsible can be located. Where a model, vendor, institution or clinician all contribute to an outcome, it may be difficult to attribute liability, and the resulting uncertainty is itself a barrier to safe AI adoption.

Some territories have rules on contributory and comparative negligence, which have been applied in medical cases. However, providing greater clarity on how these apply to AI would support adoption. In other contexts, liability standards are missing. A WHO Europe study showed that only 8 percent of countries in the region have liability standards, while 86 percent cite legal uncertainty as a major barrier to adoption.[59] This so-called 'AI liability gap' can lead clinicians to either over-defer to an AI tool or refuse to use it.[60]

National bodies that overlay AI onto existing regulation should consider the implications of designating AI tools as products or services. In common law, software has typically been treated as a service for liability purposes. This position, in which each actor is held to a standard of reasonable care, is contested and evolving.[61] However, the EU Product Liability Directive treats software in the same way as physical medical devices, placing strict liability for defective products on manufacturers and other economic operators in the supply chain.[62,63]

To ensure confident use of AI, systems should provide clear internal allocation of responsibility documented at the point of adoption. Where humans are overseeing AI outputs, they should have genuine authority, and individuals should be given the autonomy to use their own expertise to weigh outputs. This is supported by the public, with a recent Health Foundation report highlighting that 70 percent of respondents would be willing to wait longer for results if that meant a human checked the AI outputs.[64] However, this can also lead to information gating (where priority information can restrict access to other content) and congested service flow, especially where the AI system operates autonomously (such as in triage), performs as well as or better than a human, and where the developer has taken on liability for any harm.

BOX 3. DEDICATED REGULATORY TEST SANDBOXES

The UK's AI Airlock is one of the most mature sandbox systems for AI as a medical device system. Now backed by a further three years of funding (announced in April 2026), AI Airlock has provided critical insights into how AI systems can improve risk management through grounded responses (based on verifiable facts) and the implications of technologies such as AI cancer diagnostics, large language models (LLMs) and voice tools.

The Medicines and Healthcare products Regulatory Agency (MHRA) is following this system with a recently announced controlled testing environment to work with regulators to explore how AI tools could be used to better understand medicines' safety and predict potential side effects. Other countries such as Spain have also developed sandboxes for high-risk health systems. Singapore's framework exempts low-risk Software as a Medical Device (SaMD) developed by public healthcare institutions from product registration, and prioritizes safety through clinician oversight, ISO 13485 compliance, and mandatory reporting.

India's Ayushman Bharat Digital Mission (ABDM) Sandbox enables developers to test AI-enabled digital health applications in a controlled environment against national interoperability standards. This allows evidence to be generated on performance, interoperability, security and responsible data governance before systems are deployed at scale, reducing implementation risks while supporting innovation.

Source: Leal, Kohan, and Yap (2025);[65] MHRA (2026);[66] Regulations.ai (2026).[67]

4.2 BIAS, EQUITY AND SUBGROUP PERFORMANCE

Trust can only be ascertained if principles of equity are embedded through design – that is, how a tool performs for the groups it serves. Risk-prediction algorithms that target extra care for patients with complex needs, and that use predicted healthcare costs as a proxy for health need, have been shown to perpetuate racial bias. Historically, unequal access to care means that the US spends less caring for Black patients, and so the algorithm predicting healthcare costs inferred that Black people were healthier than other patients.[68] Coded by humans, the model's definition – and not its performance – was the source of inequity; only subgroup-level evaluation could expose it. A requirement for reporting, including equity impact assessments, before AI adoption is one of the most consequential and least burdensome safeguards available.

4.3 TRANSPARENCY AND INTELLIGIBILITY

Users and patients need to know, in terms they can act on, whether an AI system is being used in their care, what the system does, its reasoning or sources it relies on, how well it performs, and its limitations. Transparency does not mean disclosing proprietary internals or training data but accessible communication of intended use, performance and uncertainty.

Qure.ai is a WHO-recommended autonomous system that combines chest x-rays and a computer-aided detection tool for tuberculosis (TB) screening. Qure.AI embeds trust through its use across hundreds of health facilities and informal provider networks in India. It demonstrates its impact through collaborations with respected organizations such as the India Health Fund.[69] These partnerships support the scaling of AI when it demonstrates effectiveness, and also make the innovation more accessible, including to LMICs.[70]

4.4 DATA PRIVACY AND CYBERSECURITY

BOX 4. DATA HEADLINES

- Two-thirds of people support the use of patient data to train AI tools – but most favor conditions: consent, anonymization, limiting private access.[71]
- Structural weaknesses have exposed the healthcare data of tens of millions of people to breaches.[72]
- The average cost of a healthcare data breach has reached $7.42 million – the highest of any industry.[73]

AI introduces new cybersecurity risks for health systems. AI-enabled attacks increase the financial and operational costs of incident response and recovery. Ungoverned use by healthcare staff of generative AI tools may expand data exposure and breach impacts. Cyberattackers are also increasingly using AI to accelerate reconnaissance, identify vulnerabilities and target AI systems through techniques such as data poisoning and model manipulation.[74]

Organizations should establish and regularly test incident response protocols for AI-related failures, cybersecurity events and patient safety concerns. Systems should have clear responsibilities, escalation pathways and recovery procedures. It is also important to reduce opportunities for lateral movement through network segmentation to isolate sensitive repositories and mitigate the impact of a data breach. This can be supported through the use of blockchain technologies and AI tools to proactively patch vulnerabilities or detect threats in real time.[75]

Equity, transparency and robust cybersecurity are key to establishing and maintaining public and clinician trust, and creating a foundation for the successful and sustained deployment of AI in health systems.

SECTION 5. DEPLOYMENT

A technically capable model introduced into a health system without the necessary organizational and technical considerations may fail to deliver value. It can also disrupt clinical workflows and safety processes and undermine trust.

The necessary standards of an AI system must be met before moving towards broader considerations of operational readiness, implementation, reimbursement and life cycle assurance. Irrespective of the specific application, a set of system-level conditions are the first post-regulatory steps to support safe, effective and sustainable deployment of AI in healthcare.

5.1 A CLEAR VALUE CASE

Value is realized only when a tool fits the context. Need is defined at the locality and may differ across healthcare organizations within the same geography. It is also a dynamic construct. A clear value case expressed in clinical, operational and financial terms can inform the transition from pilot to scale. For instance, in many health systems, ambient AI scribes have demonstrated reduced clinician administrative burden, increased visit volume, and measurable reduction of self-reported burnout – from early pilots to full integration into clinical infrastructure.[76–78] But successful integration can only occur when an organization articulates the specific decision supported by the AI, the action it triggers and the measurable outcome it delivers, in advance of procuring the health intervention.

5.2 DATA QUALITY, INTEROPERABILITY AND REPRESENTATIVENESS

AI inherits the strengths and flaws of the data it learns from and the data it meets in production. While deployers may not interrogate the nuances of an approved model's development, they should be aware of the data properties that underpin the relevant outcome. For instance:

1. **Quality:** How accurate, complete, timely and consistently coded were the data used in the AI solution?
2. **Interoperability:** To what extent were the data moved across systems without loss of meaning – avoiding degradation of a model trained in one context as it is applied elsewhere?
3. **Representativeness:** To what extent do the data used for AI development represent the context it is being deployed in – to avoid inequalities becoming entrenched among under-represented subgroups?

Safe deployment therefore needs:

- An understanding of the model's reasoning to reach the relevant conclusion, or characteristics of the data used to create the model.
- How this differs from local data, population characteristics or processes.
- An articulation of this alignment for users to be aware of.

This information needs to be explicit, and purchasers should be able to access this information when making decisions.

5.3 CULTURAL RESPONSIVENESS AND SOVEREIGNTY

AI systems cannot simply be translated across contexts – they must also be adapted to the cultural, clinical and regulatory settings where they are deployed. Healthcare decisions are shaped by local epidemiology, care pathways, health-seeking behaviors, languages and societal values. A model that performs well in one setting may not transfer safely to another. This is reinforced by emerging evidence of substantial cross-lingual performance disparities. One evaluation found that OpenAI's GPT-3.5 model produced about 18 percent fewer comprehensive and appropriate responses in non-English languages and was almost 6 times more likely to generate contradictory or incorrect answers when not responding in English. This highlights the risks of introducing globally trained models without robust local evaluation.[79] While this performance can improve, the current result strengthens the case for AI sovereignty, with countries having the capability to evaluate, adapt, govern and, where appropriate, develop AI systems using local data, languages and priorities.

Therefore, AI sovereignty is not solely a question of technological independence, but of ensuring that AI reflects local needs, protects equity, and earns public trust through culturally responsive deployment.[80,81] Sovereignty also spans compute and institutional capabilities. AI readiness involves having the infrastructure to deploy, govern, evaluate and (if necessary) exit AI systems without dependence on external providers.

In the long term, this may involve building foundational models incorporating localized norms or adapting pre-existing tools to suit a particular context. In the interim, LMICs in particular will need practical frameworks for adopting cloud and generative AI on terms they can operationally control, as well as defend politically.[82]

5.4 WORKFORCE READINESS AND AI LITERACY

The workforce implications of the use of AI tools cannot be overstated; from shifting pathways to interpretation of outputs, all require the engagement of staff. The EU Artificial Intelligence Act now makes a form of AI literacy among deployers a legal obligation in the region. But as direct-to-consumer health AI increases in popularity, broader considerations are needed across the human–AI interface, encompassing hospital systems, their staff and the patients they care for. The Coalition for Health AI (CHAI) quantified this disconnect with its recent US patient survey, where 75 percent of respondents reported using AI, but only 13 percent said they feel very comfortable with it.[83] With increasing consumer-facing AI in health systems, more understanding is needed of the potential consequences of how non-AI-literate populations interact with these tools.

One of the most contentious workforce-related areas of health AI surrounds automation. In 2022, the National Health Commission of the People's Republic of China released a policy to prohibit AI substituting for physicians – reaffirming that most systems view AI as an adjunct to, not a replacement for, human involvement.[84] When well implemented, the tools showing most maturity – such as scribes, or smart scheduling services which support the prediction of non-attendance

at appointments – return time and attention to clinicians. However, these benefits depend on workflow fit, and claims of utility can only be measured in practice and not assumed.[85]

As systems evolve, some tasks will be redistributed, new responsibilities – such as oversight, validation, escalation – will be created, and some human skills will need active protection against erosion through over-reliance on AI. The OECD's workforce analysis stresses that human oversight is what makes AI principles operative, and that engagement with users is a precondition of trust and adoption.[86] Therefore, workforce planning – not only technology procurement – is part of any serious deployment strategy. If engaged early, the healthcare provider workforce can champion implementation, problem-solve deployment and use peer-support networks to act as enablers to scale AI models through knowledge mobilization.

5.5 MODERNIZED PROCUREMENT

Procurement can play a critical role in enabling the safe and scalable deployment of AI in healthcare. In countries without single-payer mechanics, modernized procurement processes that leverage centralized frameworks, certification mechanisms and standardized model evaluations can reduce the burden on individual providers. As AI technologies become more widely adopted and complex, collective procurement requirements should establish clear expectations (interoperability, data governance, cybersecurity, transparency, performance), reducing resource duplication. By embedding these standards into purchasing decisions, policymakers can create consistent incentives for innovation while ensuring that AI tools meet the safety, trust and accountability requirements of publicly funded health systems.[87]

5.6 REIMBURSEMENT AND VALUE APPRECIATION

Health AI tools are currently assessed largely through existing Health Technology Assessments (HTAs). While these frameworks remain important, they do not fully capture the key facets of AI: (i) model evolution over time; (ii) benefits (such as workflow efficiencies) realized at the health-system level; and (iii) contributions to clinical decision-making that are difficult to assign to discrete service codes.

Addressing these challenges may require changes to the reimbursement frameworks that allow developers and providers to be compensated for using AI services. The American Medical Association has proposed a new coding framework, Clinically Meaningful Algorithmic Analyses (CMAA), within the Current Procedural Terminology system.[11] CMAA codes are intended to describe services where AI analyzes clinically relevant data and generates actionable outputs, including scenarios where no physician is involved.

Alternative reimbursement pathways may also be required (see Box 5). Managed access agreements (previously used in pharmaceuticals) could support conditional entry of AI tools while additional evidence is generated. Such mechanisms support real-world evidence generation, reduce uncertainty and increase transparency. However, the path from temporary funding to long-term reimbursement remains underdeveloped.[88]

BOX 5. ADVANCING CHRONIC CARE WITH EFFECTIVE, SCALABLE SOLUTIONS (ACCESS) NATIONAL REIMBURSEMENT PROGRAM

ACCESS is one of the first national reimbursement programs specifically designed to support technology-enabled care through outcome-aligned payments. Organizations can use AI, remote monitoring, wearables, apps, software and virtual care to achieve outcomes in one of four clinical tracks, as long as the care remains clinician-guided and meets regulatory requirements.

Payments for AI-enabled clinical services are reimbursed on the basis of real-world outcomes rather than billable activities. This shows how reimbursement can support safe and sustainable scaling of AI systems in healthcare, and expands the scope for clinicians to offer innovative AI-enabled care through a streamlined reimbursement pathway.

Source: Centers for Medicare & Medicaid Services (2026).[89]

Novel reimbursement frameworks have also been proposed, linking payment levels for AI-supported diagnostic tools to the complexity of diagnosis and AI's contribution. These frameworks distinguish between assistive, augmentative and autonomous models, aligning reimbursement with the level of clinical responsibility and human oversight required. By translating these factors into a point-based reimbursement structure, the frameworks provide a transparent and auditable basis for payment, support payer confidence, and can link economic incentives with principles such as equity.[90]

An environment for effective governance is established through careful AI deployment that appreciates local context, ensures that people have the necessary skills, and that is facilitated by appropriate procurement and reimbursement.

SECTION 6. GOVERNANCE

Effective governance is vital if the health benefits of AI are to be realized. Governance allows for, and benefits from, appropriate evaluation. However, if governance responsibility falls between organizations, regulators and levels of government, gaps in oversight can slow the safe deployment of AI. To avoid this, governance must have clear lines of accountability and be proportionate to the risk involved.

6.1 UNDERSTANDING THE LIMITS OF AUTHORIZATION

A 2024 analysis of 521 FDA-authorized AI devices found that about 43 percent had no published clinical-validation data, only roughly 28 percent were validated prospectively, and just 4.2 percent had been evaluated in randomized controlled trials.[91]

A separate 2025 review of more than 1,000 device summaries found that nearly half reported no clinical study and more than half reported no performance metric, with only modest improvement after the 2021 Good Machine Learning Practice (GMLP) guidance.[92] The presence of market authorization on a regulator's list is not always a guarantee of real-world clinical effectiveness. For an organization integrating AI, one of the stages must be regulation of continuous learning models.

6.2 ORGANIZATIONAL GOVERNANCE STRUCTURES

Health organizations must define clear governance structures:

- People must be involved in the AI workflow (humans-in-the-loop).
- AI systems must be traceable to support governance oversight.
- A cross-functional AI body – across clinical, IT, cybersecurity and business – should support procurement and oversee probabilistic or dynamically learning software.
- Individuals must be accountable for each system in use.
- Operating procedures for clinicians or patients to raise concerns must be defined.
- Governance decisions must be documented.

These local processes can then contribute to regulatory system-level oversight and support robust PMS measures. For example, enterprise AI governance teams implemented by major academic health systems such as Mayo Clinic Platform or Mass General Brigham AI provide support across the AI life cycle, from use-case prioritization and procurement through deployment, monitoring and retirement.[93,94] These processes have enabled health systems to move beyond isolated pilots towards scaled deployment of AI tools by embedding interdisciplinary expertise within a single pathway, while maintaining oversight of safety, performance, equity and compliance.

6.3 RISK AT ALL LEVELS

Risk assessment is the foundation for many regulatory pathways – articulating how an AI tool can negatively impact a patient, process or pathway based on a quantitative and qualitative evaluation. While risk begins with model developers who determine the intentions and instructions for use, it is also important for teams implementing AI solutions to consider. This includes mapping how the particular solution is expected to perform with local structures and capabilities, and within the adjacent population. Tools such as the FAIR-AI Framework aim to provide a process for formal review and consistent risk assessment to support local evaluation (see Box 6).

BOX 6. FAIR-AI FRAMEWORK

The FAIR-AI Framework is a prescriptive and scalable approach to allow health leaders to evaluate AI tools. It has been used in Atrium Health in North Carolina, US. The framework begins with screening questions and, if necessary, in-depth review to assign a risk category to the proposed AI solution. The framework links this risk evaluation to an 'AI safe plan' and labeling, offering a defensible and documented decision by leaders on a particular use case. The authors suggest that FAIR-AI has categorized about 50 percent of solutions at their institution as 'low-risk', suggesting it can facilitate early adoption, ensuring that the correct solutions are implemented within a given context.

Source: Wells et al. (2025).[95]

6.4 GUARDRAILS AGAINST VALUE EGRESS

When AI deployments are scaled, who receives the value of the data? When a public system contributes its data, its clinical expertise and its patients to improve a third-party model and receives in return only a license fee and a dependency, value leaks out of the system that generated it. Placing guardrails on AI can ensure that the value accrues to the system and population, not just the vendor. Practicable measures include robust data governance, benefit-sharing agreements (rights to performance data), local evaluation, interoperable procurement standards, and arrangements to ensure that the learning generated by AI supports domestic capability.[96,97]

This is especially important for LMIC systems, which risk becoming sources of data and revenue for tools they cannot themselves shape, afford or sustain. In the global health context with cross-border threats such as antimicrobial resistance, where data is provided to multinational AI systems, it is unclear whether value or impact is realized within countries especially if resources are constrained.[98,99] Guardrails are needed to maintain data sovereignty, and to realize local value.

When such governance is lacking, the potential for disparities increases. This is already being seen in the AI trial landscape, where there are differences between high-income countries and LMICs in the quantity of AI trials, and also in the use of the AI systems.[100] The result is AI systems skewed toward non-communicable diseases, without adequate real-world evidence to support safe deployment in LMICs.[101,102]

SECTION 7. EVALUATION

Standardized evaluation generates the evidence required to demonstrate safety, effectiveness and value, supporting the shift from institutional adoption to diffusion across diverse healthcare systems. Yet, evidence suggests that only a small minority of healthcare AI pilots achieve large-scale implementation. OECD analysis found that only about 10 percent of AI imaging applications have reached national-scale deployment,[38] while the Dutch AI Coalition reported that only 3 percent of healthcare AI pilots scaled successfully.[103]

BOX 7. DEMAND FORECASTING IN SIERRA LEONE

In collaboration with the Sierra Leone Government, teams from the University of Pennsylvania have developed and implemented a novel machine learning (ML) decision support tool. The tool is decision-aware and uses multitask learning to assist in the distribution of essential medicines. It works by forecasting demand and using predictions to support allocation decision-making – leading to a 19 percent increase in consumption in districts deploying the tool.

By working in-country, the team were able to elicit real-world evidence of the utility of the AI. This has facilitated nationwide scale-up, at a cost of $30 per month covering two million women and children. The ability to readily transition from pilot to implementation highlights the value of safeguards that are flexible while not neglecting their duty. This can be a challenge in some higher-income countries where regulatory burden can stifle scale or where outcomes of real-world evaluation are not acted on at the system level. This may offer the opportunity for low- and middle-income countries to avoid the legacy systems of other countries.

Source: Chung et al. (2026).[104]

7.1 POST-MARKET SURVEILLANCE AND THE REGULATORY PATHWAY

Continuous learning capabilities require ongoing monitoring to convert one-time authorization into ongoing assurance. Lighter-touch market-entry pathways include the FDA's 510(k) pre-market route, a submission made to the FDA that demonstrates a device to be marketed is as safe, effective and "substantially equivalent" to an already legally marketed device (a predicate).[105] This pathway is faster than the De Novo application, which evaluates low-to-moderate-risk medical devices that do not have a legally marketed predicate device, and requires submitters to substantiate their equivalence claims.

However, the pathway risks 'grandfathering', where old regulations continue to be followed and, if one device is recalled, subsequent devices using this predicate are not. With AI, the pathway has allowed rapid clearances of tools, but there are concerns over 'predicate creep', with successive iteratively changing algorithms.

For instance, in October 2025, an AI stereotactic navigation system for ear, nose and throat surgery originally cleared through the 510(k) route had to be recalled after certain component revision combinations meant it could fail to meet its specified accuracy.[106]

In such cases, robust post-market surveillance is a necessity, with incident reporting for AI-related harms and near-misses seen as a routine safety function, similar to pharmacovigilance (where risks are monitored once an approved product is in use).[107] Rather than creating a stand-alone reporting mechanism for every new AI application, health systems can extend existing patient safety reporting infrastructure to capture AI-related harms and near misses, such as the UK's Learn from Patient Safety Events Service.[108] However, clinicians and health systems need to be empowered, and must understand the specific AI-related risks. The OECD's *Towards a Common Reporting Framework for AI Incidents* offers a potential foundation for healthcare AI surveillance. The framework distinguishes between AI incidents that result in harm and AI hazards that could plausibly have caused harm but were prevented.[50]

This approach mirrors established patient-safety systems and would enable clinicians to report adverse events and near misses involving AI systems. Given the speed at which dynamic algorithms can cause potential harm, there is also a need to share these events with policymakers across and between health systems. The OECD's AI Incidents and Hazards Monitor helps to share this knowledge, but it is not health specific.[109] HealthAI's Global Regulatory Network involving governments and regulatory authorities also shares knowledge among its members and provides a global platform for locally validated AI solutions and a Global Early Warning System.[110]

7.2 MAINTAINING SAFETY AT SCALE: PERFORMANCE DRIFT

Performance decay (or drift) – the gradual decline in machine learning's predictive accuracy over time – can be challenging to detect. It can occur silently as case-mix, equipment, coding practice or disease prevalence change. Testing models that assess cardiac surgery risk, acute kidney injury and radiology have all shown degradation because of differences in data distributions due to spatial or temporal variation.[111,112]

Predefined drift thresholds are essential, with transparent plans of action (recalibration, retraining with a change-control plan or withdrawal) as part of a continuous monitoring protocol. However, while some regulators have strengthened developer monitoring obligations, there is often a lack of local capability to act on them. The formation of cross-functional, interdisciplinary teams within organizations can support this capacity and ensure that any action is timely.

BOX 8. CRITICAL PERFORMANCE METRICS

While there is no consensus about critical metrics, deployers should consider:

- **Clinical performance:** Sensitivity, specificity, predictive values, precision, recall, and data distribution measures – tracked continuously.
- **Equity and fairness:** Disaggregated performance on sex, gender, ethnicity, geography, deprivation.
- **Workflow and adoption:** Use, alert overrides, clinical acceptance, user satisfaction, time-to-act.
- **Safety:** Adverse events, near misses, inappropriate recommendations, downtime, system failures, model drift (data, phenotype, concept), incidents.
- **Value:** Productivity, waiting times, length of stay, costs, throughput, patient-related outcome measures and those specific to LLMs (such as completeness, tone).

Sources: Gilbert et al. (2025);[113] Royal College of Radiologists (2026).[114]

Evaluation of AI in health systems is never complete. Because new AI often builds on previous AI applications, there needs to be continuous evaluation and consideration of the potential for drift into unevidenced domains.

SECTION 8. PUTTING EDGE AI INTO PRACTICE IN QATAR

This report's central argument is that the value of health AI is determined less by the technology and more by the readiness of the health system to receive it. Over the past decade, Qatar has deliberately invested in the foundations aligned with the EDGE AI framework to ensure that AI delivers beyond the pilot stage. This makes Qatar a useful test case, and this section looks at EDGE AI principles and how they work in practice in Qatar. Qatar demonstrates that successful AI adoption depends on building complementary capabilities across trust, deployment, governance and evaluation, rather than investing in individual technologies alone.

8.1 ESTABLISHING TRUST

Qatar has been building trust across all its AI capabilities, starting with improving support of its data assets. For instance, Qatar's genomics consent and bioethics infrastructure, set alongside its data-protection regime, provides a basis for public confidence (see Box 9). The country's investment in a representative national data asset is itself a hedge against the leakage of value which smaller systems are exposed to. The National Skilling Program, delivered in conjunction with Microsoft, has trained 50,000 people in advanced digital skills, complementing capacity-building across universities and hospitals.[115] This has helped improve widespread use of AI and digital tools, and the understanding of their value. Individual programs such as breast cancer screening have improved clinician confidence in AI through interactive workshops, peer learning with clinicians from neighboring countries, comparison of AI outputs and review of real-world cases.

BOX 9. THE QATAR GENOME PROGRAM AND THE QATAR PRECISION HEALTH INSTITUTE

Qatar's most distinctive asset lies in population genomics. Qatar Biobank (2012) and the Qatar Genome Program (2015) gave Qatar the region's first population-based genome program. Since 2024 the two have been combined under the Qatar Precision Health Institute and have recruited about 50,000 participants and sequenced tens of thousands of whole genomes from Qatari and wider Arab populations. This now acts as reference data for a long-standing under-represented group in global databases. Combined with the public system's shared electronic health record, this creates a highly integrated base for precision medicine and AI tools. Early clinical translation, such as the return of actionable findings to participants on BRCA1/2 (the human genes that suppress tumors and repair DNA), illustrates the opportunity provided by shared data, but also the careful governance that responsible deployment demands.

Sources: Mbarek et al. (2022);[116] QPHI (no date).[117]

Qatar is also advancing trust by advocating for international co-ordination on AI governance through the launch of the Global Alliance for AI Ethics (GAAIE). This multistakeholder initiative intends to strengthen global dialogue on the ethical development and use of AI. The Alliance aims to ensure that diverse ethical perspectives and cultural traditions inform the development of international technical standards and regulatory frameworks.[118]

BOX 10. QATAR COMPUTING RESEARCH INSTITUTE (QCRI) AT HAMAD BIN KHALIFA UNIVERSITY

The QCRI's Qatar Center for Artificial Intelligence pursues AI for precision health. One of the Center's most notable outputs has been **Fanar**, an Arabic LLM, reflecting a deliberate emphasis on linguistic and cultural representativeness. Qatar's National AI Strategy requires that AI respect local language and culture. Its sovereign Arabic-language capability positions it to develop and evaluate clinical AI aligned to the languages and norms of the population it serves.

Source: QCRI (2026).[119]

8.2 DEPLOYMENT

Care in Qatar is delivered principally through three public institutions: Hamad Medical Corporation (HMC), the main provider of secondary and tertiary care; the Primary Health Care Corporation (PHCC), which operates a primary-care network of more than 30 health centers; and Sidra Medicine, Qatar Foundation's specialist women's and children's hospital and research center.

The significant structural advantage that Qatar offers to AI deployment is data continuity. HMC and PHCC have operated a shared electronic health record since 2012, under a program that is among the first initiatives in the region to place a public health system on a single platform.[120] This interoperable, consistently coded data is the foundation on which safe and generalizable AI deployment ultimately depends.

As a result, concrete clinical deployments are emerging, with real-world patient benefits. These include the adoption of AI-enabled adaptive radiotherapy by the National Center for Cancer Care and Research, and an oral-cancer detection tool (DenTech) developed jointly by the Qatar Computing Research Institute (QCRI) and Qatar University collaborating with HMC.[121]

8.3 GOVERNANCE

Qatar's robust governance system to facilitate its health ambitions is outlined in the National Vision 2030; the Third National Development Strategy; and the National Health Strategy 2024–2030, which aims for a fully digital, integrated and preventive system.[122] A national AI policy runs in parallel, as set out in the 2019 National AI Strategy.

A cross-government AI Committee was established under Cabinet Decision No. (10) of 2021 to co-ordinate adoption across sectors.[123] The wider digital transformation framework is provided by the Digital Agenda 2030, with healthcare among the priority sectors for implementation.[124]

Cross-cutting safeguards include the Personal Data Privacy Protection Law, and guidance from the National Cyber Security Agency (NCSA) on secure AI adoption[125] to ensure that robust oversight forms a part of implementation.

Qatar's AI guidelines also help close governance gaps for AI systems that fall outside medical device regulation. The guidelines require non-clinical AI applications to demonstrate fairness, accountability and transparency through documented intended use, development processes, known limitations, compliance with data protection laws, and other governance requirements. This provides a proportionate oversight framework for AI systems that may influence health services or public outcomes, despite not meeting the definition of a regulated medical device.[126]

8.4 EVALUATION

In AI terms, 'explainability' means the ability to describe how choices have been made in clear human terms; and a 'black box problem' is when AI fails to achieve this. Qatar advocates for explainability as a key means of addressing black box problems post-implementation. The country recommends classification thresholds and provides requirements for explaining AI decision-making processes. These components contribute to Qatar's post-market monitoring expectations, such as ensuring system reliability, enabling contestability, and providing avenues for redress. This sets up a basis for continuous monitoring, as well as empowering those in authority to act on the outputs from evaluation.[126,127]

BOX 11. AI-ASSISTED MAMMOGRAPHY

Interpretation systems in the National Breast Cancer Screening program highlight how Qatar is adapting to the need for continuous oversight, audit and performance review of deployed AI systems. The AI tool, which integrates into screening workflows, provides automated image analysis and risk scoring for screening mammograms. Responsibility for final interpretation and reporting remains with radiologists – ensuring human oversight with clear accountability. Prior to deployment, the program undertook retrospective evaluation and validation activities to assess the performance of the technology within the local screening environment. Initial challenges were due to the Primary Health Care Corporation's (PHCC's) wider IT systems upgrade and contract renewal process, which slightly delayed implementation but enabled a more robust pilot with relevant teams.

> "If implementing the program again we would place even greater emphasis on early governance, expectation-setting, stakeholder engagement, local validation, and post-deployment monitoring."
> – PHCC program representative.

Following deployment, PHCC established ongoing monitoring and evaluation processes to assess program performance and support continuous improvement, providing a template for life cycle assurance.

The evaluation has shown that the proportion of women recalled for further assessment decreased from 20.4 percent to 15.5 percent, demonstrating meaningful real-world improvements in clinical true positive rate (31.6 percent to 38.3 percent), and a reduction in the false positive rate (18.4 percent to 14.1 percent).

Source: BioSpectrum (2024);[128] PHCC (2026).[129]

Having established many of the EDGE AI foundations for safe AI, Qatar now faces the challenge common to mature health systems: demonstrating real-world value and sustaining performance as AI scales. The country's experience shows that AI readiness is built progressively through investment in health-system capabilities suited to national context and priorities.

SECTION 9. RECOMMENDATIONS

AI is already delivering value across health systems in diagnostics, clinical administration, patient navigation and service planning. However, the conditions required for safe, equitable and sustainable deployment remain uneven. As outlined in Section 3, the governance, evidence and operational requirements for AI should be proportionate to the nature of the system, the level of risk it poses, and the risk tolerance of the context where AI is used.

The recommendations in this section are intended to help policymakers establish the conditions necessary for effective AI deployment. While some apply across all AI systems, most are designed to be implemented using a risk-calibrated approach, ensuring that oversight is proportionate to potential harms. AI systems should also be interpreted in the context of a country's existing level of AI maturity, recognizing that implementation will differ between health systems with well-established digital infrastructure and governance arrangements and those where these capabilities are still emerging.

9.1 ESTABLISHING TRUST

Recommendation 1. Equity by design

Foundational action: Mandate subgroup performance reporting (by age, sex, ethnicity and other clinically relevant populations) alongside a pre-deployment equity impact assessment demonstrating local validation and, where necessary, adaptation to the user population.

Advanced action: Make disaggregated equity reporting a condition of national adoption or reimbursement. Require reporting throughout rather than as a one-off submission, with reporting intervals determined by the risk of the AI system and the stability of its performance. Higher-risk or rapidly evolving systems should undergo more frequent review, with reporting intervals extended only once sustained performance and equitable outcomes have been demonstrated.

Recommendation 2. Deliberate human–AI interaction and trust

Foundational action: For each AI use case, specify what information is presented, to whom, and when, with a working override mechanism. This can contain automation bias and alert fatigue.

Advanced action: Embed participation of patients/patient representatives in oversight committees, evaluation processes, and post-deployment review, to ensure that patient perspectives inform decisions on safety, equity, transparency and acceptable risk.

9.2 DEPLOYMENT

Recommendation 3. Govern the deployment, not just the product (pre-scale gate)

Foundational action: Before scaling any AI system, establish implementation readiness through appropriate digital infrastructure, interoperable workflows and local governance. Require a prospective local validation study, with predefined sample sizes and pass/fail thresholds in the intended deployment population, distinct from the original regulatory evaluation, with a mandatory 'go/no-go' gate to control information flow for high-risk systems before wider deployment.

Advanced action: Make risk-tiered minimum performance thresholds (sensitivity, specificity and calibration) a condition of deployment approval. These thresholds should be defined at the point of authorization and serve as the minimum standard a system must meet before scaling. They should be revisited only if the model, intended use or deployment context changes. This assessment determines whether a system is suitable for deployment at scale; ongoing performance after deployment should instead be governed through continuous post-market monitoring, as outlined in Recommendation 7.

Recommendation 4. Procurement for outcomes, interoperability and adaptability

Foundational action: Reimburse AI systems based on demonstrated real-world clinical or operational outcomes in the deploying health system, reviewed at fixed intervals (eg, every 12 to 24 months), rather than based solely on pre-market accuracy claims or regulatory clearance. As a condition of procurement, there should be a requirement for adherence to recognized standards, data portability and contractual safeguards against becoming dependent on a single supplier.

Advanced action: There should be a requirement for cross-functional procurement panels (clinical, technical, financial) to assess AI technologies against common risk-based evaluation criteria. These panels should contribute to a formal buyer network to facilitate sharing validation results, implementation experience and purchasing information to avoid information asymmetry and duplication.

9.3 GOVERNANCE

Recommendation 5. Patient-facing tools governed by function and risk

Foundational action: Close the oversight gap for general-purpose AI tools, including foundation models, by introducing a risk-calibrated registration process. Any system that extends beyond general wellness information to influence healthcare decisions – including decisions about whether, when or where to seek care – should be registered and subject to proportionate regulatory oversight.

Advanced action: Limit reimbursement to tools built for specific health uses that define the use cases and limitations, sourcing standards, and specific safeguards for defined vulnerable groups (such as children or patients at risk of mental illness).

Recommendation 6. Named accountability

Foundational action: Fund workforce AI-skill training and role redesign at the point of deployment. Before going live, put accountability in writing across institution, vendor and clinician.

Advanced action: Designate institutional AI governance capacity within each provider with the formal authority to suspend or withdraw a system.

9.4 EVALUATION

Recommendation 7. Resource monitoring with thresholds and set deadlines (post-deployment life cycle)

Foundational action: Once live, require ongoing monitoring with regular risk-calibrated reporting (eg, quarterly), with a named accountable owner and a scheduled formal review (eg, annually), independent of whether performance has changed.

Advanced action: Fund continuous automated performance monitoring against predefined drift thresholds that trigger proportionate action, such as recalibration, retraining or suspension, when exceeded. Complement this with a mandatory expiration clause requiring periodic re-authorization at risk-proportionate intervals (eg, up to 5 years for lower-risk systems) to confirm continued safety, effectiveness and fitness for purpose. These drift thresholds should remain distinct from the fixed minimum performance thresholds established before deployment (Recommendation 3).

Recommendation 8. Independent evaluation capability

Foundational action: Fund national or pooled regional assurance functions to validate tools against local data, run sandboxes, set benchmarks and establish post-market surveillance – pooling across smaller systems where scale alone will not support this.

Advanced action: Establish a low-friction national AI registry to capture approvals, deployments, real-world performance and safety incidents, enabling shared learning across institutions and contributing to international AI incident reporting networks.

9.5 CONCLUSION: A CALL TO ACTION

The decisions made over the next two to three years will determine how AI is integrated into healthcare for the decade ahead. Governments, regulators, funders and providers are now establishing the governance, procurement, interoperability and digital sovereignty arrangements that will shape whether AI delivers meaningful improvements in healthcare quality, innovation and system performance.

The impact of AI in health will depend less on the sophistication of algorithms than on the strength of the health systems that deploy them. Health systems need the institutional, digital and workforce capabilities described in the EDGE AI framework to deploy, govern and continuously evaluate AI while maintaining trust, safety and equity.

Countries that invest in these capabilities will be better positioned to adopt AI safely and equitably while strengthening health system resilience and sustainability. Policymakers should move beyond enabling innovation alone and focus on building the governance, evidence and implementation capacity to enable AI to deliver measurable improvements in health outcomes at scale.

APPENDIX. KEY AI SOURCE DOCUMENTS

MODEL TESTING/PRE-MARKET VALIDATION

Territory/Body	Key documents – linked	Published (and updated)	Summary of key considerations
European Union	Document 32024R1689 AI in healthcare (Commission) + Medical Devices Regulation (MDR)/ In Vitro Diagnostic Regulation (IVDR)	AI Act in force Aug 2024; MDR 2017/745 & IVDR 2017/746	Most clinical AI is classed as high-risk, triggering a dual conformity assessment under the AI Act and the MDR/IVDR (a single notified-body route is allowed). Pre-market obligations: a risk-management system; high-quality, representative and governed training/validation/ test data sets; technical documentation; defined accuracy, robustness and cybersecurity; transparency; and human oversight designed in. Binding and legally enforceable.
United Kingdom	Medicines and Healthcare products Regulatory Agency (MHRA): Software & AI as a Medical Device Change Programme (roadmap)	Program from 2021; guidance updated Feb 2025; AI Airlock sandbox pilot completed Mar 2025	Risk-based life cycle approach layered on the UK medical device regime. Focus on defining intended purpose and classification, data quality and bias mitigation, demonstrating real-world performance, and managing adaptive/continuous-learning models. The AI Airlock – the first regulatory sandbox for AI as a medical device – lets manufacturers test novel products directly with the regulator. Designed to align with National Institute for Health and Care Excellence (NICE) and National Health Service (NHS) routes.
United States	US Food and Drug Administration (FDA) AI-Enabled Device Software Functions – lifecycle Management and Marketing Submission Recommendations: Draft Guidance Good Machine Learning Practice for Medical Device Development: Guiding principles	GMLP Oct 2021; total product life cycle (TPLC) guidance issued in draft 7 Jan 2025	TPLC approach overlaid on existing pre-market pathways (510(k), De Novo, premarket approval (PMA). Submissions are expected to cover model description, data lineage/representativeness and train-test splits, performance tied to the intended claims, bias analysis and mitigation, the human–AI workflow and human factors, transparency (eg, model cards) and cybersecurity. Good Machine Learning Practice (GMLP) (jointly with Health Canada and the MHRA) sets cross-cutting good-practice principles. The life cycle guidance is still a draft.
China	National Medical Products Administration (NMPA): Classification & Definition of AI Medical Software (No. 47, 2021) Center for Medical Device Evaluation of NMPA: Registration Review of AI Medical Devices (No. 8, 2022)	Classification Jul 2021; registration-review guidelines Mar 2022	Risk-classified registration with whole-life cycle management. Emphasis on data quality, algorithm verification and validation, clinical evaluation, and review of source code and the quality-management system. A special approval (innovation) channel accelerates review for qualifying devices, supported by real-world-data pilots; the Personal Information Protection Law governs health-data use. Binding.

Territory/Body	Key documents – linked	Published (and updated)	Summary of key considerations
Singapore	Health Sciences Authority (HSA): Digital Health: Understanding digital health products and the regulations Ministry of Health (MOH)/HSA: Emerging regulatory policy issues: AI in Healthcare Guidelines (AIHGle 2.0)	Software as a Medical Device (SaMD) life cycle guidelines 2019 (updated); AIHGle Oct 2021, AIHGle 2.0 (10 Mar 2026)	Life cycle, risk-based SaMD regulation by HSA, complemented by the MOH/HSA AIHGle good-practice guidance for developers, deployers and users – covering data quality, validation, transparency and clinical evaluation. AIHGle 2.0 gives worked examples of when an AI tool is a regulated medical device (AI-SaMD). An international reliance pathway leverages approvals from reference regulators to speed access.
Kenya	Pharmacy and Poisons Board (PPB) Guideline on Regulation of Medical Device Software (MDSW) in Kenya – Software as a Medical Device/ Software in a Medical Device (SaMD/SiMD) Artificial Intelligence Bill 2026 (tabled) Digital Health Act 2023 (No. 15 of 2023)	Digital Health Act 2023; National AI Strategy Mar 2025; PPB MDSW guideline 2026; AI Bill 2026 tabled in the Senate (under deliberation)	The PPB regulates Medical Device Software – both SaMD and SiMD – through a risk-based, International Medical Device Regulators Forum (IMDRF)-aligned framework requiring pre-market approval, with submissions covering software design documentation, risk assessment, clinical validation and version control, and cybersecurity built in. AI-enabled tools face a higher bar: developers must show that training data is broad and representative across patient profiles and the system was properly tested. The draft AI Bill 2026 would classify healthcare AI as high-risk, requiring pre-deployment risk assessment and human oversight.
Organisation for Economic Co-operation and Development (OECD)	OECD AI Principles (Recommendation on AI) Scaling Artificial Intelligence in Health	AI Principles 2019, updated May 2024; health report Mar 2026	Non-binding intergovernmental standard. Pre-deployment expectations are framed through robustness, security and safety, transparency and explainability, and accountability across the life cycle. The health report's policy checklist stresses data foundations, validation and assurance before scaling. The OECD definition of an AI system is reused by the EU, US and others, aiding interoperability.
World Health Organization (WHO)	Regulatory considerations on AI for health Ethics and governance of AI for health: Guidance on Large Multi-Modal Models (LMMs)	Regulatory considerations 19 Oct 2023; LMM guidance 18 Jan 2024	Non-binding guidance for regulators and developers. Eighteen considerations across six areas: documentation and transparency of development and validation; clearly stated intended purpose; data set selection and reference standards; analytical and clinical validation; total life cycle risk management; and external validation across diverse populations to mitigate bias. The LMM guidance extends validation thinking to generative models.

POST-MARKET SURVEILLANCE AND MONITORING

Territory/Body	Key documents – linked	Published (and updated)	Summary of key considerations
European Union	Document 32024R1689 – EU AI Act – post-market monitoring & incident reporting Artificial Intelligence in Healthcare	AI Act Post-market Surveillance (PMS) duties for high-risk areas apply from Dec 2027 (Aug 2028 for AI in medical devices) under the 2026 Digital Omnibus; MDR/IVDR vigilance ongoing	Providers must operate a post-market monitoring system, log performance and report serious incidents and malfunctions. MDR/IVDR add vigilance, post-market clinical follow-up and periodic safety update reporting. Significant changes or retraining can trigger re-assessment, and deployers must monitor real-world use and maintain human oversight.
United Kingdom	Medical devices: post-market surveillance regime (in force) Software and AI as a Medical Device: Change Programme roadmap – PMS workstream	PMS statutory instrument signed Dec 2024; in force 16 Jun 2025	Strengthened statutory PMS duties: trend and periodic safety reporting, faster (serious) incident reporting and field-safety corrective actions, with specific attention to performance drift in adaptive AI. Surveillance findings contribute to the wider Software and AI as a Medical Device change program.
United States	FDA final: Predetermined Change Control Plan (PCCP) for AI-Enabled DSFs Predetermined Change Control Plans for Machine Learning-Enabled Medical Devices: Guiding principles	PCCP final guidance Dec 2024; life cycle monitoring (draft) Jan 2025	A PCCP lets a manufacturer pre-authorize a defined envelope of future modifications – with protocols for data, retraining, testing, labeling and deployment – without a new submission, provided that the intended use is unchanged. The life cycle guidance asks for post-market performance monitoring plans; mandatory adverse event (MDR/MedWatch) reporting and recalls continue to apply.
China	NMPA registration-review guidelines (change control) NMPA adverse-event monitoring and real-world data programs	Guidelines 2022; real-world-data pilots ongoing	Whole-life cycle management requires change control and re-registration for significant algorithm updates, alongside adverse event reporting. Real-world data programs are used to support ongoing evaluation and re-approval. Oversight is centralized (126 Class III AI devices approved by Dec 2024).
Singapore	MOH Emerging regulatory policy issues: AIHGle 2.0 – deployment & monitoring guidance HSA Digital Health: Understanding digital health products and the regulations	AIHGle 2.0 Mar 2026; HSA PMS ongoing	Deployers and users are expected to assess and mitigate deployment risks, monitor real-world performance and model drift, keep humans in the loop and maintain feedback channels. HSA field safety and adverse event reporting applies to AI-SaMD, with explicit attention to generative and continuously learning AI.

Territory/Body	Key documents – linked	Published (and updated)	Summary of key considerations
Kenya	PPB Guideline on Regulation of Medical Device Software in Kenya (MDSW) AI Bill 2026 (registration & auditing of high-risk systems)	Act 2023; PPB guideline and AI Bill 2026	PPB approval requires manufacturers to commit to ongoing performance monitoring and post-market surveillance after release, with cybersecurity safeguards such as encryption and access controls. The draft AI Bill would subject high-risk healthcare systems to mandatory registration and continuous auditing, and place obligations on both providers and deployers. Enforcement is still being built out.
OECD	OECD Principles Scaling AI in Health (guardrails)	Principles 2019/2024; health report Mar 2026	Recommends ongoing monitoring, incident reporting and measurement of safety, privacy and fairness after deployment. The health checklist's 'guardrails' pillar covers oversight and monitoring, and points to a developing AI incident reporting framework and a global regulatory network for incident sharing. Non-binding.
WHO	Regulatory considerations on AI for health (post-deployment monitoring) Ethics and governance of AI for health: Guidance on LMMs	2023/2024	Advocates total life cycle risk management, continuous performance monitoring, and post-deployment surveillance for safety, bias and drift, with collaboration across regulators. The LMM guidance specifically flags monitoring for misinformation, fabrication ('hallucination') and inequitable performance. Non-binding.

DISSEMINATION AND COMMISSIONING

Territory/Body	Key documents – linked	Published (and updated)	Summary of key considerations
European Union	European Health Data Space (EHDS) Regulation Document 32024R1689 – AI Act deployer duties + Health Technology Assessment (HTA) Regulation	EHDS adopted 2025 (phased); HTA Regulation joint clinical assessments phased in from Jan 2025 (initially cancer medicines and ATMPs)	The EHDS enables secondary use of health data and cross-border infrastructure for developing and validating AI. The HTA Regulation introduces a joint EU clinical assessment relevant to adoption decisions. Deployers (hospitals, GP practices) carry obligations they cannot offload to vendors. Commissioning itself remains largely national.
United Kingdom	National Institute for Health and Care Excellence (NICE): Evidence Standards Framework (ESF) for Digital Health Technologies	First 2019; updated Apr 2021 and Aug 2022 (adds AI/adaptive algorithms + deployment standards)	The ESF (not mandatory) sets evidence tiers for clinical performance and economic value that National Health Service (NHS) commissioners and purchasers use to decide whether to adopt a technology. It spans design, value, performance, delivery and deployment, and requires a plan to monitor performance and retraining over time. The NHS AI in Health and Care Award funded real-world evaluation (program now closed), and NICE's Office for Digital Health works with the Medicines and Healthcare products Regulatory Agency (MHRA) on access pathways.

Territory/Body	Key documents – linked	Published (and updated)	Summary of key considerations
United States	FDA PCCP – user communication of changes Good Machine Learning Practice for Medical Device Development: Guiding principles	Ongoing; voluntary assurance/certification efforts maturing 2025	There is no central commissioning framework. Adoption is market- and payer-driven, via Centers for Medicare & Medicaid Services and private-payer coverage and procedure codes rather than national commissioning. The FDA requires authorized changes to be communicated to users, and voluntary assurance bodies (eg, Coalition for Health AI (CHAI), working with The Joint Commission) are emerging as quasi-standards for trust.
China	NMPA approval + New-Generation AI Development Plan	AI Development Plan 2017; procurement/reimbursement policy ongoing	State-led diffusion: central approval combines with national strategy, public procurement and hospital pricing/reimbursement to drive scale, with a strong push into imaging and triage. Data-localization and Personal Information Protection Law shapes what can be deployed and shared. Adoption is fast where products fit national priorities.
Singapore	MOH Emerging regulatory policy issues: AIHGle 2.0 – deployment & monitoring guidance HSA Digital Health: Understanding digital health products and the regulations	AIHGle 2.0 Mar 2026; reliance pathway from 2018	National co-ordination through MOH, HSA and the health-tech agency. AIHGle clarifies the distinct accountability of developers, deployers and users, and stresses transparency to support trusted deployment, while the reliance pathway speeds market access. Institutional readiness and governance are emphasized over volume.
Kenya	The Digital Health (Health Information Management Procedures) Regulations National AI Strategy 2025–2030	Act 2023; Regulations Apr 2025; Strategy launched 27 Mar 2025	State co-ordinated adoption. The Digital Health Agency certifies digital health solutions against a certification framework and mandates interoperability via the Kenya Health Enterprise Architecture. The National AI Strategy sets adoption, infrastructure and capacity priorities and flags health use cases. There is no evidence-for-commissioning instrument; diffusion runs through certification, interoperability standards and public investment.
OECD	Scaling AI in Health (engagement + trustworthy deployment) OECD AI Principles	Health report Mar 2026; Principles 2019/2024	The policy checklist frames adoption around data foundations, capacity-building, meaningful engagement with the public, providers and industry, and assured trustworthy deployment. It highlights workforce training, reimbursement clarity and procurement, and is built around comparative learning across member states. Non-binding.
WHO	Regulatory considerations on AI for health (collaboration & data ecosystems) Ethics and governance of AI for health: Guidance on LMMs	2023 / 2024 (Global Strategy on Digital Health 2020–2025)	Encourages stakeholder collaboration, shared data ecosystems and equitable access, with governments urged to provide public or not-for-profit infrastructure in exchange for adherence to ethical principles. Strong emphasis on dissemination that reaches low-resource settings rather than widening inequities. Non-binding.

ABBREVIATIONS

ABDM	Ayushman Bharat Digital Mission
ACCESS	Advancing Chronic Care with Effective, Scalable Solutions
AI	artificial intelligence
AIDRS	AI and Digital Regulations Service
AIHGle	Artificial Intelligence in Healthcare Guidelines
CHAI	Coalition for Health AI
CMAA	Clinically Meaningful Algorithmic Analyses
EHDS	European Health Data Space
EU	European Union
FDA	United States Food and Drug Administration
GAAIE	Global Alliance for AI Ethics
GMLP	Good Machine Learning Practice
HSA	Health Sciences Authority (Singapore)
HMC	Hamad Medical Corporation
HTA	Health Technology Assessment
IMDRF	International Medical Device Regulators Forum
IVDR	In Vitro Diagnostic Regulation
LLM	large language model
LMICs	low- and middle-income countries
LMM	large multi-modal models
MDR	Medical Devices Regulation
MDSW	Medical Device Software
MHRA	Medicines and Healthcare products Regulatory Agency (UK)
MOH	Ministry of Health (Singapore)
MoPH	Ministry of Public Health (Qatar)
ML	machine learning
NCSA	National Cyber Security Agency
NHS	National Health Service (UK)
NICE	National Institute for Health and Care Excellence (UK)

OECD	Organisation for Economic Co-operation and Development
PCCP	Predetermined Change Control Plan
PHCC	Primary Health Care Corporation
PMS	post-market surveillance
PPB	Pharmacy and Poisons Board (Kenya)
QCRI	Qatar Computing Research Institute
SaMD	Software as a Medical Device
SiMD	Software in a Medical Device
TB	tuberculosis
TPLC	total product life cycle
US	United States of America
WHO	World Health Organization

ACKNOWLEDGMENTS

The Forum Advisory Group for this paper was co-chaired by **Professor Lord Ara Darzi**, Director of the Institute of Global Health Innovation, Imperial College London, and **Professor Ahmed Elmagarmid**, Executive Director of the Qatar Computing Research Institute. The authors acknowledge the support received from the Qatar Foundation in creating this report, and the input received from the tech industry.

Sincere thanks are extended to the members of the WISH 2026 Forum on **Artificial Intelligence Expert Advisory Group**, who contributed their unique insights and collective experience to this paper:

- **Professor Alastair Denniston**, Professor of Regulatory Science and Innovation, Consultant Ophthalmologist, University of Birmingham
- **Dr Dominic King**, Vice President of Health, Microsoft AI
- **Professor Kee Yuan Ngiam**, Professor of Surgery, Head, Artificial Intelligence Office, National University Health System, Singapore
- **Associate Professor Ziad Obermeyer**, Blue Cross of California Distinguished Associate Professor, UC Berkeley School of Public Health
- **Andy Song**, Director and Special Counsel, Google DeepMind
- **Professor Effy Vayena**, Swiss Federal Institute of Technology Zurich (ETH Zurich)

The co-chairs and authors would also like to extend their sincere thanks to the excellent WISH8 team: **Dr Slim Slama**, **Dr Taskeen Khan** and **Sangeeta Singh**, as well as **Dominic Harris** for their editorial guidance and contributions.

REFERENCES

1. US Food and Drug Administration (FDA). Artificial Intelligence-Enabled Medical Devices. FDA [Preprint]; 2026a. https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-enabled-medical-devices (accessed 15 July 2026).
2. Wang J et al. AI-Enhanced Handheld ECGs to Screen for Prior Myocardial Infarction in Rural India. Working Paper Series. Cambridge, MA: National Bureau of Economic Research; 2026. https://doi.org/10.3386/w34690.
3. Niu C and Wang G. The current and future landscape of AI foundation models for cancer management. Nature Communications, 17(1), p. 5277; 2026. https://doi.org/10.1038/s41467-026-72576-5.
4. Costa-Gomes B et al. Public use of a generalist LLM chatbot for health queries. Nature Health, 1(7), pp. 689–696; 2026. https://doi.org/10.1038/s44360-026-00117-x.
5. Wong A et al. External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients. JAMA Internal Medicine, 181(8), pp. 1065–1070; 2021. https://doi.org/10.1001/jamainternmed.2021.2626.
6. Bean AM et al. Reliability of LLMs as medical assistants for the general public: A randomized preregistered study. Nature Medicine, 32(2), pp. 609–615; 2026. https://doi.org/10.1038/s41591-025-04074-y.
7. Hussein R et al. Advancing healthcare AI governance through a comprehensive maturity model based on systematic review. npj Digital Medicine, 9(1), p. 236; 2026. https://doi.org/10.1038/s41746-026-02418-7.
8. European Union. Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC (2024). 2024a. http://data.europa.eu/eli/dir/2024/2853/oj (accessed 16 July 2026).
9. European Union. High-level summary of the AI Act | EU Artificial Intelligence Act. 2024b. https://artificialintelligenceact.eu/high-level-summary (accessed 15 July 2026).
10. Medicines and Healthcare products Regulatory Agency (MHRA). Impact of AI on the regulation of medical products: Implementing the AI White Paper principles. London: MHRA; 2024a. https://assets.publishing.service.gov.uk/media/662fce1e9e82181baa98a988/MHRA_Impact-of-AI-on-the-regulation-of-medical-products.pdf.
11. American Medical Association (AMA). CPT Appendix S: Taxonomy for artificial intelligence in medical services & procedures. 2026. https://www.ama-assn.org/practice-management/cpt/cpt-appendix-s-taxonomy-artificial-intelligence-medical-services-procedures (accessed 15 July 2026).

12. Tsinghua University. Tsinghua University holds Tsinghua AI Agent Hospital Inauguration and 2025 Tsinghua Medicine Townhall Meeting. 2025. https://www.tsinghua.edu.cn/en/info/1245/14224.htm (accessed 15 July 2026).

13. Philips. Future Health Index: AI in practice Shaping the future of healthcare now. Philips; 2026. https://www.philips.com/c-dam/corporate/newscenter/global/pdf/philips-future-health-index-2026-report-ai-in-practice.pdf.

14. World Health Organization (WHO). Health workforce. 2025b. https://www.who.int/health-topics/health-workforce (accessed 15 July 2026).

15. Epic. Penny for revenue cycle and operations. 2026. https://www.epic.com/software/penny (accessed 15 July 2026).

16. One Advanced. GP Workflow Optimisation through AI-Assisted workflows. 2026. https://www.oneadvanced.com/products/gp-workflow-assistant (accessed 15 July 2026).

17. Singhal K et al. Large language models encode clinical knowledge. Nature, 620, pp. 172–180; 2023. 2023. https://www.nature.com/articles/s41586-023-06291-2.

18. Shah M et al. Towards conversational medical AI with eyes, ears and a voice. arXiv, 2026. https://doi.org/10.48550/arXiv.2605.09272.

19. Feng X et al. The Effectiveness of AI Chatbots in Alleviating Mental Distress and Promoting Health Behaviors Among Adolescents and Young Adults: Systematic Review and Meta-Analysis. Journal of Medical Internet Research, 27(1), p. e79850; 2025. https://doi.org/10.2196/79850.

20. Li D et al. Use of Artificial Intelligence-Assisted Conversational Agents to Improve Patient Experience Related to Physicians: Cross-Sectional Study in China. Journal of Medical Internet Research, 27, p. e76540; 2025. https://doi.org/10.2196/76540.

21. Nature Health. Consumer-facing AI can improve public health – but more evidence is needed. Nature Health, 1(2), pp. 159–159; 2026. https://doi.org/10.1038/s44360-025-00052-3.

22. Menz BD et al. Current safeguards, risk mitigation, and transparency measures of large language models against the generation of health disinformation: repeated cross sectional analysis. BMJ, 384, p. e078538; 2024. https://doi.org/10.1136/bmj-2023-078538.

23. US Food and Drug Administration (FDA). Examples of Software Functions for Which the FDA Will Exercise Enforcement Discretion. FDA [Preprint]; 2025a. https://www.fda.gov/medical-devices/device-software-functions-including-mobile-medical-applications/examples-software-functions-which-fda-will-exercise-enforcement-discretion (accessed 15 July 2026).

24. Makary MA and Prasad V. Priorities for a New FDA. JAMA, 334(7), p. 565; 2025. https://doi.org/10.1001/jama.2025.10116.

25. Duffourc MN, Goel R and Kennedy R. Regulation of Consumer Safety Risks from AI Characters in the European Union and United States. In: Crea C and Infantino M, editors. Remedies to Digital Vulnerability in European Private Law. Cham: Springer Nature Switzerland, pp. 59–82; 2025. https://doi.org/10.1007/978-3-032-19720-7_4.

26. California Legislative Information. California Legislative Information – Chapter 22.6. Companion Chatbots. 2026. https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?lawCode=BPC&division=8.&title=&part=&chapter=22.6.&article= (accessed 15 July 2026).

27. US Food and Drug Administration (FDA). General Wellness: Policy for Low Risk Devices Guidance for Industry and Food and Drug Administration Staff. 2026b. https://www.fda.gov/media/90652/download.

28. Reuters. Harvard Medical School licenses consumer health content to Microsoft. 9 October 2025. https://www.reuters.com/business/healthcare-pharmaceuticals/microsoft-partner-with-harvard-healthcare-push-cut-openai-reliance-wsj-reports-2025-10-08 (accessed 15 July 2026).

29. Gross B et al. Introducing Copilot Health. Microsoft AI; 2026. https://microsoft.ai/story/introducing-copilot-health (accessed 15 July 2026).

30. Ferber D et al. Towards autonomous medical artificial intelligence agents. Nature, pp. 1–10; 2026. https://doi.org/10.1038/s41586-026-10675-5.

31. Collaco BG et al. The role of agentic artificial intelligence in healthcare: A scoping review. npj Digital Medicine, 9(1), p. 345; 2026. https://doi.org/10.1038/s41746-026-02517-5.

32. Wang Q et al. A feasibility study of automating radiotherapy planning with large language model agents. Physics in Medicine & Biology, 70(7), p. 075007; 2025. https://doi.org/10.1088/1361-6560/adbff1.

33. Srinivasu PN et al. Exploring agentic AI in healthcare: A study on its working mechanism. Frontiers in Medicine, 12, p. 1753443; January 2026. https://doi.org/10.3389/fmed.2025.1753443.

34. Moss W et al. Mapping AI regulation in health care with the Health & AI Policy Index. npj Digital Medicine, 9(1), p. 413; 2026. https://doi.org/10.1038/s41746-026-02734-y.

35. Artificial Intelligence/Machine Learning-enabled Working Group (IMDRF). Good machine learning practice for medical device development: Guiding principles. International Medical Device Regulators Forum; 2025. https://www.imdrf.org/sites/default/files/2025-02/IMDRF_AIML%20WG_GMLP_N88%20Final.pdf.

36. World Health Organization (WHO). Regulatory considerations on artificial intelligence for health. Geneva: WHO; 2023. https://www.who.int/publications/i/item/9789240078871 (accessed 15 July 2026).

37. World Health Organization (WHO). Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models. Geneva: WHO; 2025a. https://www.who.int/publications/i/item/9789240084759.

38. Organisation for Economic Co-operation and Development (OECD). Scaling Artificial Intelligence in Health. Paris: OECD Publishing; 2026. https://doi.org/10.1787/a436e12d-en.

39. Foundation for Law and International Affairs. China's New Generation of Artificial Intelligence Development Plan. 2017. https://flia.org/notice-state-council-issuing-new-generation-artificial-intelligence-development-plan (accessed 15 July 2026).

40. National Medical Products Administration (NMPA). NMPA announcement on guidance for the classification defining of AI-based medical software products. 2021. https://english.nmpa.gov.cn/2021-07/08/c_660267.htm (accessed 15 July 2026).

41. European Commission. Artificial Intelligence Act. 2024. http://data.europa.eu/eli/reg/2024/1689/oj.

42. European Commission. Regulation (EU) 2025/327 on the European Health Data Space and amending Directive 2011/24/EU and Regulation (EU) 2024/2847 – Public Health. 2025. https://health.ec.europa.eu/publications/regulation-eu-2025327-european-health-data-space-and-amending-directive-201124eu-and-regulation-eu_en (accessed 15 July 2026).

43. Han Y, Ceross A and Bergmann J. Regulatory frameworks for AI-enabled medical device software in China: Comparative analysis and review of implications for global manufacturer. JMIR AI, 3, p. e46871; 2024. Available at: https://doi.org/10.2196/46871.

44. Pharmacy and Poisons Board. Guideline on Regulation of Medical Device Software in Kenya (MDSW). 2026. https://web.pharmacyboardkenya.org/download/guideline-on-regulation-of-medical-device-software-in-kenya-mdsw (accessed 15 July 2026).

45. International Medical Device Regulators Forum (IMDRF). Characterization considerations for medical device software and software-specific risk. IMDRF; 2025. https://www.imdrf.org/documents/characterization-considerations-medical-device-software-and-software-specific-risk (accessed 15 July 2026).

46. Medicines and Healthcare products Regulatory Agency (MHRA). Software and AI as a medical device change programme roadmap. GOV.UK; 2024b. https://www.gov.uk/government/publications/software-and-ai-as-a-medical-device-change-programme/software-and-ai-as-a-medical-device-change-programme-roadmap (accessed 15 July 2026).

47. Medicines and Healthcare products Regulatory Agency (MHRA). Software and artificial intelligence (AI) as a medical device. GOV.UK; 2025. https://www.gov.uk/government/publications/software-and-artificial-intelligence-ai-as-a-medical-device/software-and-artificial-intelligence-ai-as-a-medical-device (accessed 15 July 2026).

48. US Food and Drug Administration (FDA). Predetermined change control plans for machine learning-enabled medical devices: Guiding principles. FDA [Preprint]; 2025b. https://www.fda.gov/medical-devices/

software-medical-device-samd/predetermined-change-control-plans-machine-learning-enabled-medical-devices-guiding-principles (accessed 15 July 2026).

49. Ministry of Health Singapore. Artificial Intelligence in Healthcare Guidelines. Version 2.0. Ministry of Health Singapore and Health Sciences Authority; 2026. https://isomer-user-content.by.gov.sg/3/6636a140-7986-43f6-9861-096333138d59/AIHGle_2.0.pdf.

50. Organisation for Economic Co-operation and Development (OECD). Towards a common reporting framework for AI incidents. OECD Artificial Intelligence Papers [Preprint]; 2025. https://doi.org/10.1787/f326d4ac-en.

51. Health Canada. Guiding principles: Predetermined change control plans for machine learning-enabled medical devices. 2024. https://www.canada.ca/en/health-canada/services/drugs-health-products/medical-devices/good-machine-learning-practice-medical-device-development/predetermined-change-control-plans-machine-learning-enabled-medical-devices.html (accessed 30 July 2026).

52. Liu Y, Yu W and Dillon T. Regulatory responses and approval status of artificial intelligence medical devices with a focus on China. npj Digital Medicine, 7(1), p. 255; 2024. https://doi.org/10.1038/s41746-024-01254-x.

53. Rahimzadeh V. US regulation of medical artificial intelligence and machine learning (AI/ML) research and development. In: Solaiman B and Cohen IG, editors. Research Handbook on Health, AI and the Law. Cheltenham, UK: Edward Elgar Publishing Ltd; 2024. http://www.ncbi.nlm.nih.gov/books/NBK613220 (accessed 16 July 2026).

54. Directorate-General for Health and Food Safety (European Commission) et al. Study on the deployment of AI in healthcare: Final report. Publications Office of the European Union; 2025. https://data.europa.eu/doi/10.2875/2169577.

55. National Health Service (NHS). AI and Digital Regulations Service for Health and Social Care. No date. https://digitalregulations.innovation.nhs.uk (accessed 16 July 2026).

56. Health Sciences Authority (HSA). Regulatory reliance. 2026. https://www.hsa.gov.sg/international-collaboration/md/regulatory-reliance (accessed 16 July 2026).

57. Arboit L et al. Surgeons' awareness, expectations, and involvement with artificial intelligence: a survey pre and post the GPT era. European Journal of Surgical Oncology, 51(12), p. 110525; 2025. https://doi.org/10.1016/j.ejso.2025.110525.

58. Nouis SC, Uren V and Jariwala S. Evaluating accountability, transparency, and bias in AI-assisted healthcare decision-making: A qualitative study of healthcare professionals' perspectives in the UK. BMC Medical Ethics, 26, p. 89; 2025. https://doi.org/10.1186/s12910-025-01243-z.

59. World Health Organization (WHO). Is your doctor's AI safe? 19 November 2025. 2025c. https://www.who.int/europe/news/item/19-11-2025-is-your-doctor-s-ai-safe (accessed 16 July 2026).

60. Medical Protection. Closing the AI liability gap: AI, safety, and the case for legislative change. 2026. https://www.medicalprotection.org/docs/mp/0/media-and-policy/policy/closing-the-ai-liability-gap-medical-protection-final.pdf.

61. Kuhl CB. Social media cases JCCP no. 5255. Westlaw, Los Angeles County Superior Court; 2023.

62. European Commission. Report on the safety and liability implications of artificial intelligence, the Internet of Things and robotics. 2020. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52020DC0064.

63. European Parliament. Directive (EU) 2024/2853 of the European Parliament. 2024. https://eur-lex.europa.eu/eli/dir/2024/2853/oj/eng.

64. Binesmael A et al. Attitudes to technology and AI in health care. The Health Foundation; 2026. https://www.health.org.uk/reports-and-analysis/analysis/attitudes-to-technology-and-ai-in-health-care (accessed 16 July 2026).

65. Leal A, Kohan PE, and Yap P. AI governance in health; 2025. https://healthai.agency/app/uploads/2025/12/AI-GOVERNANCE-IN-HEALTH-REPORT-2025_v2.pdf.

66. Medicines & Healthcare products Regulatory Agency (MHRA). MHRA launches AI sandbox to accelerate medicines development and improve safety. GOV.UK; 2026. https://www.gov.uk/government/news/mhra-launches-ai-sandbox-to-accelerate-medicines-development-and-improve-safety (accessed 30 July 2026).

67. Regulations.ai. Spain – National AI Sandbox (Royal Decree 817/2023): National AI Sandbox. 2026. https://regulations.ai/regulations/RAI-ES-NA-NASFSXX-2025 (accessed 30 July 2026).

68. Obermeyer Z et al. Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), pp. 447–453; 2019. https://doi.org/10.1126/science.aax2342.

69. Vijayan S et al. Artificial intelligence-assisted chest X-ray for tuberculosis case finding in low- and middle-income countries: Implementation experiences and impact. BJRIOpen, 8(1), p. tzag007; 2026. https://doi.org/10.1093/bjro/tzag007.

70. Qure.ai. A tale of two cities: AI for health equity. 6 September 2023. https://www.qure.ai/impact-stories/a-tale-of-two%20cities-AI-for-health-equity (accessed 16 July 2026).

71. Public First. Public First startup coalition polling. Polling tables 24 June 2026. 2026. https://www.publicfirst.co.uk/public_first_startup_coalition_polling (accessed 16 July 2026).

72. Freyer O et al. Are we heading towards a cybersecurity crisis in health care and are actions needed? The Lancet Digital Health, 8(1); 2026. https://doi.org/10.1016/j.landig.2025.100946.

73. IBM. Cost of a Data Breach Report 2025: The AI Oversight Gap. IBM; 2025. https://www-api.ibm.com/adobe/assets/urn:aaid:aem:607b9590-38e0-4c91-b433-aa8a17f5b5e8/original/as/cost-of-a-data-breach-2025-full-report.pdf.

74. Brooks C. Why AI Is Creating New Cybersecurity Risks for Healthcare. Forbes; 2026. https://www.forbes.com/sites/chuckbrooks/2026/06/09/ai-cybersecurity-risks-healthcare (accessed 16 July 2026).

75. Di Palma G et al. AI-induced cybersecurity risks in healthcare: A narrative review of blockchain-based solutions within a clinical risk management framework. Risk Management and Healthcare Policy, 18, pp. 3479–3497; 2025. https://doi.org/10.2147/RMHP.S544523.

76. Lukac PJ et al. A randomized-clinical trial of two ambient artificial intelligence scribes: Measuring documentation efficiency and physician burnout. medRxiv: The Preprint Server for Health Sciences; 2025. https://doi.org/10.1101/2025.07.10.25331333.

77. Olson KD, Meeker D and Troup M. Use of ambient AI scribes to reduce administrative burden and professional burnout. JAMA Network Open, 8(10), p. e2534976; 2025. https://doi.org/10.1001/jamanetworkopen.2025.34976.

78. Holmgren AJ et al. Ambient artificial intelligence scribes and physician financial productivity. JAMA Network Open, 9(1), p. e2553233; 2026. https://doi.org/10.1001/jamanetworkopen.2025.53233.

79. Jin Y et al. Better to ask in English: Cross-lingual evaluation of large language models for healthcare queries. Proceedings of the ACM Web Conference 2024. New York, NY, USA: Association for Computing Machinery (WWW '24), pp. 2627–2638; 2024. https://doi.org/10.1145/3589334.3645643.

80. Mügge D. EU AI sovereignty: For whom, to what end, and to whose benefit? Journal of European Public Policy, 31(8), pp. 2200–2225; 2024. https://doi.org/10.1080/13501763.2024.2318475.

81. Independent Expert Group on Artificial Intelligence and Culture. Report of the Independent Expert Group on Artificial Intelligence and Culture. UNESCO; 2025. https://www.unesco.org/sites/default/files/medias/fichiers/2025/09/CULTAI_Report%20of%20the%20Independent%20Expert%20Group%20on%20Artificial%20Intelligence%20and%20Culture%20%28final%20online%20version%29%201.pdf.

82. Vota W. Compute reality of artificial intelligence in global health LMICs. ICTworks, 28 May 2026. https://www.ictworks.org/compute-reality-of-ai-in-global-health-lmics (accessed 30 July 2026).

83. Coalition for Health AI (CHAI). Patient survey report on health AI & transparency. CHAI; 2026. https://www.chai.org/blog/chai-releases-new-patient-survey-report-on-health-ai-and-transparency (accessed 16 July 2026).

84. National Health Commission of the People's Republic of China, Office Medical Affairs. Notice on issuing the detailed rules for the supervision of internet-based medical services. 2022. https://www.nhc.gov.cn/yzygj/c100068/202203/2072f0e8988249e59d942e1b2a933916.shtml (accessed 16 July 2026).

85. Dai T, McDonald KM and Baumgart DC. Global advances in health artificial intelligence: A workforce imperative. The Lancet, 0(0); 2026. https://doi.org/10.1016/S0140-6736(26)00693-8.

86. Almyranti M et al. Artificial intelligence and the health workforce: Perspectives from medical associations on AI in health. OECD Artificial Intelligence Papers [Preprint]; 2024. https://doi.org/10.1787/9a31d8af-en.

87. Evans TD et al. The role of procurement frameworks in responsible AI innovation in the National Health Service: A multi-stakeholder perspective. Frontiers in Health Services, 5, p. 1608087; 2025. https://doi.org/10.3389/frhs.2025.1608087.

88. Nelson Advisors. The French PECAN framework at the three year milestone. Nelson Advisors Blog; 2026. https://www.healthcare.digital/single-post/the-french-pecan-framework-at-the-three-year-milestone (accessed 16 July 2026).

89. Centers for Medicare and Medicaid Services. ACCESS (Advancing Chronic Care with Effective, Scalable Solutions) Model. 2026. https://www.cms.gov/priorities/innovation/innovation-models/access (accessed 16 July 2026).

90. Kirchhoff J et al. Valuing diagnostic AI: A structured reimbursement model for learning healthcare systems. Frontiers in Digital Health, 7; 2025. https://doi.org/10.3389/fdgth.2025.1642750.

91. Chouffani El Fassi S et al. Not all AI health tools with regulatory authorization are clinically validated. Nature Medicine, 30(10), pp. 2718–2720; 2024. https://doi.org/10.1038/s41591-024-03203-3.

92. Mehta V et al. Evaluating transparency in AI/ML model characteristics for FDA-reviewed medical devices. npj Digital Medicine, 8(1), p. 673; 2025. https://doi.org/10.1038/s41746-025-02052-9.

93. Yu Y et al. Accelerating AI innovation in healthcare: Real-world clinical research applications on the Mayo Clinic Platform. npj Health Systems, 3(1), p. 17; 2026. https://doi.org/10.1038/s44401-026-00068-1.

94. Mass General Brigham. Mass General Brigham AI. No date. https://www.massgeneralbrigham.org/en/research-and-innovation/centers-and-programs/artificial-intelligence/who-we-are/leadership-team (accessed 16 July 2026).

95. Wells BJ et al. A practical framework for appropriate implementation and review of artificial intelligence (FAIR-AI) in healthcare. npj Digital Medicine, 8(1), p. 514; 2025. https://doi.org/10.1038/s41746-025-01900-y.

96. Morley J et al. The ethics of AI in health care: A mapping review. Social Science & Medicine, 260, p. 113172; 2020. https://doi.org/10.1016/j.socscimed.2020.113172.

97. Meskó B and Topol EJ. The imperative for regulatory oversight of large language models (or generative AI) in healthcare. npj Digital Medicine, 6(1), p. 120; 2023. https://doi.org/10.1038/s41746-023-00873-0.

98. Qin H et al. Towards Trustworthy Artificial Intelligence for Equitable Global Health. arXiv; 2023. https://doi.org/10.48550/arXiv.2309.05088.

99. Ong JCL et al. International partnership for governing generative artificial intelligence models in medicine. Nature Medicine, 31(9), pp. 2836–2839; 2025. https://doi.org/10.1038/s41591-025-03787-4.

100. Gumede SB et al. Harnessing artificial intelligence in health research in low-income and middle-income countries: Potential and caution. Mayo Clinic Proceedings: Digital Health, 4(3), p. 100380; 2026. https://doi.org/10.1016/j.mcpdig.2026.100380.

101. Ciecierski-Holmes T et al. Artificial intelligence for strengthening healthcare systems in low- and middle-income countries: A systematic scoping review. npj Digital Medicine, 5(1), p. 162; 2022. https://doi.org/10.1038/s41746-022-00700-y.

102. Gah EE. Algorithmic fairness in clinical machine learning for non-communicable disease prediction in Sub-Saharan Africa: A systematic review, empirical analysis, and research agenda. Rochester, NY: Social Science Research Network; 2026. https://doi.org/10.2139/ssrn.6839418.

103. AI Coalition 4 NL (AIC4NL). Only 3% of AI pilots in healthcare scale up: AIC4NL seeks solutions for structural acceleration. 5 November 2025. https://aic4nl.nl/en/werkgebieden/gezondheid-en-zorg/slechts-3-van-ai-pilots-in-de-zorg-schaalt-op-aic4nl-zoekt-oplossingen-voor-structurele-versnelling (accessed 16 July 2026).

104. Chung AT-H et al. Improving access to essential medicines via decision-aware machine learning. Nature, 653(8116), pp. 1178–1183; 2026. https://doi.org/10.1038/s41586-026-10433-7.

105. US Food and Drug Administration (FDA). Premarket Notification 510(k). FDA; 2024. https://www.fda.gov/medical-devices/premarket-submissions-selecting-and-preparing-correct-submission/premarket-notification-510k (accessed 16 July 2026).

106. US Food and Drug Administration (FDA). Class 2 Device Recall TruDi Navigation System. Recall Z-0420-2026. 2026c. https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfres/res.cfm?id=216327 (accessed 15 July 2026).

107. Guni A et al. Global regulation of surgical AI devices: Assessing cross-border compatibility in the United States, Europe, United Kingdom, and China. International Journal of Surgery, Ovid [Preprint]; 2026. https://doi.org/10.1097/JS9.0000000000005378.

108. NHS England. Learn from patient safety events (LFPSE) service. No date. https://www.england.nhs.uk/patient-safety/patient-safety-insight/learning-from-patient-safety-events/learn-from-patient-safety-events-service (accessed 16 July 2026).

109. Organisation for Economic Co-operation and Development (OECD). AIM: AI Incidents and Hazards Monitor. No date. https://oecd.ai/en/incidents (accessed 16 July 2026).

110. HealthAI. Global Regulatory Network (GRN). No date. https://healthai.agency/grn_1 (accessed 30 July 2026).

111. Vela D et al. Temporal quality degradation in AI models. Scientific Reports, 12(1), p. 11654; 2022. https://doi.org/10.1038/s41598-022-15245-z.

112. Rotalinti Y et al. Identifying and understanding significant change due to drift when assessing AI models in healthcare: A narrative review. BMJ Digital Health & AI, 2(1); 2026. https://doi.org/10.1136/bmjdh-2026-000085.

113. Gilbert S et al. Could transparent model cards with layered accessible information drive trust and safety in health AI? npj Digital Medicine, 8(1), p. 124; 2025. https://doi.org/10.1038/s41746-025-01482-9.

114. Royal College of Radiologists. Post-deployment monitoring and safety reporting of AI medical imaging devices in clinical practice. London: Royal College of Radiologists; 2026. https://www.rcr.ac.uk/media/baphq10s/rcr-publications-post-deployment-monitoring-and-safety-reporting-of-ai-medical-imaging-devices-in-clinical-practice.pdf.

115. Ministry of Communications and Information Technology (MCIT). MCIT, Microsoft Launch National Skilling Program, Inaugurate Digital Center of Excellence. MCIT; 2022. https://www.mcit.gov.qa/en/news/mcit-microsoft-launch-national-skilling-program-inaugurate-digital-center-of-excellence (accessed 16 July 2026).

116. Mbarek H et al. Qatar genome: Insights on genomics from the Middle East. Human Mutation, 43(4), pp. 499–510; 2022. https://doi.org/10.1002/humu.24336.

117. Qatar Precision Health Institute (QPHI). Personalized care, healthier living. No date. https://www.qphi.org.qa (accessed 16 July 2026).

118. Ministry of Communications and Information Technology (MCIT). State of Qatar announces the launch of the Global Alliance for AI Ethics (GAAIE) during its participation in the Global Dialogue on AI Governance in Geneva. State of Qatar; 7 July 2026. https://www.mcit.gov.qa/en/news/State-of-Qatar-Announces-the-Launch-of-the-Global-Alliance-for-AI-Ethics (accessed 16 July 2026).

119. Qatar Computing Research Institute (QCRI). QCRI website. 2026. https://www.hbku.edu.qa/en/qcri (accessed 16 July 2026).

120. Hamad Medical Corporation (HMC). Integration of Oracle Cerner Electronic Health Record across Hamad Medical Corporation and Qatar Biobank helps accelerate scientific research. 2023. https://hamad.qa:443/EN/news/2023/June/Pages/Integration-of-Oracle-Cerner-Electronic-Health-Record-Across-Hamad-Medical-Corporation.aspx (accessed 16 July 2026).

121. Hamad Bin Khalifa University. DenTech. 2025. https://www.hbku.edu.qa/en/web-summit-2026/startups/dentech (accessed 16 July 2026).

122. Ministry of Public Health (MoPH). National Health Strategy. State of Qatar; 2024. https://www.moph.gov.qa/english/NHS/Pages/The-Aspiration.aspx (accessed 16 July 2026).

123. Ministry of Communications and Information Technology (MCIT). National Artificial Intelligence Strategy for Qatar 2019. State of Qatar; 2019. https://www.mcit.gov.qa/-/media/mcit/documents/strategies/national_artificial_intelligence_strategy_for_qatar_2019_en.pdf.

124. Regulations.ai. Qatar – Digital Transformation Strategy: Qatar Digital Agenda 2030. 2024. https://regulations.ai/regulations/RAI-QA-NA-QDA2XXX-2024.

125. National Cyber Security Agency (NCSA). Guidelines for Secure Adoption and Use of AI. 2024. https://regulations.ai/regulations/RAI-QA-NA-GSAUAXX-2024 (accessed 16 July 2026).

126. Solaiman B et al. A "True Lifecycle Approach" towards governing healthcare AI with the GCC as a global governance model. npj Digital Medicine, 8(1), p. 337; 2025. https://doi.org/10.1038/s41746-025-01614-1.

127. Solaiman B, Bashir A and Dieng F. Regulating AI in health in the Middle East: Case studies from Qatar, Saudi Arabia and the United Arab Emirates. In: Solaiman B and Cohen IG, editors. Research Handbook on Health, AI and the Law. Cheltenham, UK: Edward Elgar Publishing Ltd; 2024. http://www.ncbi.nlm.nih.gov/books/NBK613206 (accessed 30 July 2026).

128. BioSpectrum. First nationwide adoption of AI in national-level breast screening workflow in Middle East. 2024. https://www.biospectrumasia.com/news/105/24604/first-nationwide-adoption-of-ai-in-national-level-breast-screening-workflow-in-middle-east.html (Accessed 4 June 2026).

129. Primary Health Care Corporation. Impact of AI Integration within the National Breast Cancer Screening Program. Internal programme evaluation and audit data covering January 2023–March 2026; 2026.

ISBN 978-1-913991-45-6



www.wish.org.qa