

April 06, 2025

21st Century Healthcare, 20th Century Rules - Bridging the AI Regulation Gap

Co-Authors:

Vaishnavi Singh, Era Sarda, Christen Rao, Romano Tucci

Addressed to:

World Health Organization (WHO) and The Honorable Robert F. Kennedy Jr., U.S. Health
Official (Secretary-designate of Health and Human Services)

Table of Contents:

1. Problem Statement..... 3

2. Background..... 3

3. Recommendations & Implementation Plan..... 5

 Engaging Educational Institutes for Local Insights..... 7

4. Impact Assessment: Regulatory Reform for AI in Healthcare..... 9

5. References & Appendices..... 11

Executive Summary

The rapid integration of artificial intelligence into clinical decision-making represents both unprecedented opportunities and significant risks. Globally, AI systems are implemented increasingly to diagnose diseases, predict patient outcomes, and guide treatment protocols. Despite this, our regulatory frameworks remain dangerously antiquated and designed for an era of static medical devices rather than adaptive, learning algorithms.

The stark disparity between healthcare AI innovation and regulatory oversight constitutes an urgent public health concern. Current fragmented approaches leave critical gaps in governance, allowing AI-driven diagnostic and decision-support tools to enter clinical settings without adequate safeguards or clear accountability structures. We must establish comprehensive, dynamic oversight mechanisms that evolve alongside the technologies they govern. The evidence demonstrates that one-time approvals and static validation protocols are fundamentally insufficient for systems that continuously learn and adapt. The time for action is now, as 2025 is anticipated to be pivotal for AI validation and regulatory approaches.

We therefore in this report herein we propose a three-pillar regulatory framework:

First, nations ought to explore implementing risk-based classification systems that apply proportionate oversight based on an AI system's potential impact on patient care. High-risk applications must face more stringent monitoring requirements with mechanisms for rapid intervention when safety concerns arise.

Second, nations must eventually mandate continuous performance monitoring across healthcare institutions through automated systems that track key performance indicators, detect anomalies, and alert regulators to potential issues. This approach acknowledges that AI risks are often silent and systemic, making them particularly dangerous in healthcare contexts where patients are inherently vulnerable.

Third, establish regulatory sandboxes with strict entry criteria to enable controlled testing of emerging AI technologies before widespread deployment. These environments must balance innovation with rigorous safeguards, ensuring new systems demonstrate consistent performance across diverse populations.

Given the global nature of healthcare technology markets, we must pursue international regulatory harmonization while respecting regional resource constraints and cultural contexts.

1. Problem Statement

Artificial intelligence (AI) is rapidly becoming a co-pilot in clinical decision-making - diagnosing diseases, predicting patient risks, and guiding treatments - yet our regulatory "rulebook" is struggling to keep pace. There is a stark mismatch between the breakneck innovation of healthcare AI and the lagging, fragmented regulations meant to govern its use. Healthcare may be more heavily regulated than sectors like finance on paper, but current frameworks remain ill-suited for the AI era (Palaniappan et al., 2024). Many of today's rules were built for a time when medical devices were static and algorithms didn't adapt

themselves; they are now outdated and piecemeal, leaving dangerous gaps in oversight. As a result, AI-driven diagnostic and decision-support tools are increasingly deployed in clinics without sufficient safeguards or clarity of accountability (*WHO Outlines Considerations*, 2023). This is not a distant theoretical concern - it's a pressing patient safety issue: unclear regulation and lack of modern oversight increase the risk of biased or unsafe AI tools slipping into medical use (Grant, 2023). We face a critical problem: healthcare AI is moving fast and breaking things, and regulation must urgently catch up.

2. Background

AI is revolutionizing healthcare by enabling faster, more accurate diagnostics, patient triage, and treatment planning (Reader, 2025). The pandemic accelerated the adoption of digital health technologies, and AI systems are now widely deployed in clinics.

Healthcare AI poses distinct challenges. Unlike conventional software, these systems continuously learn and evolve based on new data. Regulatory frameworks built on one-time approvals and static validation are insufficient for such tools (Palaniappan et al., 2024). There is no infrastructure, like aviation's "black box" systems, for real-time monitoring and error detection in hospitals (Xu & Shuttleworth, 2023). As a result, AI misdiagnoses and performance drifts may go unnoticed, silently harming patients.

Adaptive algorithms add further complexity. While continuous learning is a key advantage of AI, unsupervised updates introduce risks. Freezing algorithms at the point of approval stifles innovation, but failing to regulate updates invites unpredictable failures. Although the World Health Organization (WHO) has called for dynamic certification systems and algorithm change protocols, implementation has been slow and inconsistent (*WHO Outlines Considerations*, 2023).

Fairness and bias are also pressing concerns. AI tools often reflect or amplify biases present in the training data. One notable case revealed that a hospital risk prediction algorithm systematically deprioritized Black patients due to flawed assumptions in its design (Vartan, 2024). Existing regulations focus on efficacy and safety but neglect equity. Moreover, when errors occur, it's unclear who is responsible, the developer, the provider, or the healthcare institution, highlighting a lack of accountability structures (Fotheringham & Smith, 2024).

Other industries offer useful regulatory models. In aviation and autonomous vehicles, rigorous testing, phased approvals, and constant monitoring are the norm (*How AI Is Making Autonomous Vehicles Safer* | *Stanford HAI*, n.d.). These sectors mandate human oversight and require systems to log decisions for post-incident analysis. Healthcare could benefit from similar practices, ensuring diagnostic AI tools are continuously evaluated and adjusted for safety.

Fintech provides another compelling model through the use of regulatory sandboxes (Leckenby et al., 2021), controlled environments where emerging tools are tested under supervision before wide release. Healthcare is beginning to experiment with this approach (*Medicines and Healthcare products Regulatory Agency*, 2025). Such sandboxes could enable real-world testing of AI systems while ensuring safeguards are in place. Fintech has

also embraced adaptive regulation, which evolves with the technology. Healthcare could follow suit by allowing for iterative approvals of AI tools, provided they meet ongoing safety benchmarks.

Yet, healthcare demands higher regulatory standards due to the stakes involved. Mistakes in this domain cost lives, not just time or money. Healthcare also entails unique ethical considerations such as consent, data privacy, and the principle of "do no harm." Sandboxes in healthcare must therefore have stricter entry criteria, and real-time monitoring must respect patient confidentiality. Fairness is not a luxury in this space, it is essential, and any bias in medical AI can have life-threatening consequences (*WHO Outlines Considerations*, 2023).

The urgency for reform is being acknowledged. WHO has emphasized the need for transparent, lifecycle regulation of AI, protection against bias, and international cooperation. Director-General Dr. Tedros has warned that AI can do harm if not properly governed (*WHO Issues First Global Report*, 2021). In the U.S., Robert F. Kennedy Jr. has advocated for AI to bridge healthcare gaps, particularly in rural areas (*RFK Jr.'s Health Tech Fix*, 2025). The FDA and public-private initiatives like the Coalition for Health AI are beginning to respond, building frameworks for responsible AI integration.

Despite these efforts, significant gaps remain. Healthcare AI is still largely shaped by private companies, with minimal input from bioethicists or public health experts. Success stories like AlphaFold dominate public discourse, while deeper systemic risks, like inequitable access (Global North vs Global South), biased algorithms, and opaque decision-making, go unaddressed. Insurance companies increasingly rely on AI to make decisions about care and coverage, reinforcing disparities under the guise of optimization (*AlphaFold*, 2025).

Unlike misinformation, which incites visible societal reactions, healthcare AI risks are silent, systemic, and harder to detect. AI operates quietly within trusted institutions. Patients rarely question diagnoses, lack the means for second opinions, and engage with healthcare only when already vulnerable. This makes errors harder to trace and even harder to challenge, especially for marginalized communities.

This hackathon proposal responds to these challenges by advocating for smart, proportionate regulation of healthcare AI. Solutions include dynamic sandbox testing, real-time performance auditing, and flexible, adaptive oversight systems. By designing frameworks that evolve alongside technology, we can ensure healthcare AI remains ethical, transparent, and equitable, protecting progress and people.

3. Recommendations & Implementation Plan

Policy Landscape:

Regulatory systems differ widely across countries. The United States has the FDA overseeing AI as part of Software as a Medical Device (SaMD) and has begun exploring real-time oversight for adaptive models (Pesapane et al., 2018; Schulz et al., 2019). The European Union's GDPR and new Medical Devices Regulation provide overlapping

oversight on data use and diagnostic functions; the new EU AI Act introduces high-risk classification and transparency requirements (Pesapane et al., 2018; Kalodanis et al., 2025). In developing countries we observe jurisdictions like Ghana face regulatory fragmentation and resource gaps that hamper effective AI governance; concerns include explainability, accountability, and culturally mismatched system design (Kayusi et al., 2024; Chavula et al., 2024).

Discussion on addressing AI risks through policy

Most current healthcare regulations were designed for fixed devices, not AI systems that learn and adapt over time (Bedoya et al., 2022). It is unclear how novel algorithms would impact healthcare systems further where they have access to large troves of patient data. Many have proposed systems in which a Deep Reinforcement Learning model would have access to real-time patient information in order to assist clinicians in diagnostics and other related processes.

Value-based Q-learning networks have great benefits such as improving estimate accuracy by using adaptive dynamic weight functions and leveraging human expertise in decision-making processes. These networks have demonstrated superior performance compared to both human clinicians and other value-based DRL methods, achieving high survival rates (up to 97.81% in the MIMIC-III dataset) when properly implemented. They can also be designed to maintain interpretability through feature selection algorithms like random forest (Wu et al., 2023).

However, DRL models also pose the unique risk of system malfunctions that could affect multiple interconnected healthcare processes simultaneously. DRL models operate on state-action-reward sequences and malfunctions could mean that models “gamify” the process of executing actions to attain rewards. Think of a model that helps prescribe medicine, or issues billing. Model drifting here would mean a lack of consistency in the care provided by clinicians, in which even a smote of imprecision could result in damage or death.

A policy framework that could address this emerging issue pre-emptively to the propagation of DRL systems in hospitals would be a Supervisory Technology approach (SupTech). We do think that there is room for policymakers to consider SupTech a valid pathway to overseeing the current efforts for SaMD in the United States, and other efforts in the world.

SupTech should enable continuous surveillance throughout the AI lifecycle rather than one-time approvals. A potential SupTech system that could be developed for hospital systems is one that can implement automated monitoring systems that track AI performance in real-time across healthcare institutions.

Learned global health leaders could look to establish a cloud-based platform similar to Singapore’s HEALIX to act as an AI “technology factory” that facilitates secure sharing of anonymized clinical and genomic data. This would require significant investment and collaboration to acquire such data but would be very beneficial in the event that widespread adoption of AI tools is certain.

Our SupTech proposal would see the deploying of AI systems to monitor other AI systems, creating a layered oversight approach. Given the adaptive nature of AI systems that learn and evolve, we would issue time-limited approvals that require periodic reassessment.

This proposal would require transparency from developers about algorithm changes, with mandatory updates to regulatory bodies when modifications occur, in alignment with SupTech proposals commonly seen in the finance world. We could implement automated alert systems when AI performance deviates from approved parameters, to jumpstart response measures.

To address the implementation of DRL models specifically, we can look to the Dynamic Activity-Aware Health Monitoring strategy (DActAHM) that utilizes an innovative approach to balance monitoring performance with efficiency in which this framework demonstrates how models optimize resource use while maintaining performance standards.

For regulatory evaluation specifically, the CORE-MD clinical risk score has been developed for medical devices incorporating AI or machine learning algorithms, which could potentially be adapted for continuous monitoring of DRL systems. It is up to supervisory agencies to direct how these 2 examples could potentially be integrated into a larger, all-encompassing SupTech system.

It is demonstrated that the lack of post-deployment oversight leads to real-world diagnostic errors not captured during development (Kan et al., 2024; Naili et al., 2025; Bedoya et al., 2022).

Furthermore, it is true that models can often show that performance metrics often ignore contextual and individualized patient information (Reyna et al., 2022). In the contrary situation, Deep learning algorithms can predict race with a high level of accuracy from a wide variety of radiological images.

Policies must be able to address AI's limitations in ignoring contextual information, but also be prepared for high-end models being able to accurately evince conclusions with limited information and what that would mean for privacy of patients within our hospital systems employing such models.

These models (both rudimentary or novel) are often AI algorithms commonly trained on biased or demographically narrow datasets, risking skewed or unsafe outcomes when generalized to broader populations (Kan, 2024; Naili et al., 2025; Reyna et al., 2022). Many systems also lack proper data preprocessing to mitigate input biases, and the use of synthetically generated data can further amplify these risks. While synthetic data is often promoted as a privacy-preserving solution, it can still retain identifying markers that violate privacy regulations. Generative Adversarial Networks (GANs) require real patient data as input, and the synthetic outputs may inadvertently preserve patterns that enable re-identification of individuals. This creates a harrowing scenario where synthetic data might be used as a mechanism to evade data privacy laws, since technically it "does not relate to an individual" in the typical sense and thus falls outside standard data protection legislation like the GDPR (Arora et al., 2022).

The risk of deanonymization is growing as techniques become more sophisticated such as in DRL models. We've already seen that personal information can be extracted from seemingly anonymous data sources like EEGs, retinal fundus photography, and even ECGs. What appears adequately anonymized today may be vulnerable to re-identification tomorrow, especially as data is increasingly shared in online repositories. Such a risk is not one we are prepared for adequately.

Perhaps more concerning is how synthetic data can perpetuate and amplify existing biases in healthcare. When GANs generate synthetic data based on biased input datasets, they effectively repackage these biases, be it in terms of their skewed composition or otherwise, and perpetuate it further through the data produced (Cross et al., 2024). These biases arise throughout the AI lifecycle, from the starting point of data collection and labeling to model development, evaluation, and deployment. When synthetic data is generated from these biased sources, it can create a false sense of data diversity while actually reinforcing the same problematic patterns.

From a policy perspective, we need frameworks that address both privacy and bias concerns in synthetic data generation. This requires a "privacy-by-design" mindset coupled with rigorous bias testing across different demographic groups. We should consider implementing a digital chain of custody for synthetic data to ensure integrity throughout its lifecycle, along with standardized bias reporting requirements that force developers to demonstrate equitable performance across population subgroups.

The healthcare community must collaborate with regulatory agencies to develop appropriate governance structures for synthetic data, ensuring that innovation proceeds without compromising patient privacy or exacerbating healthcare disparities. Without careful oversight, synthetic data could become a technological means of perpetuating long-standing healthcare inequities under the ill-fitted notion of innovation. (Giuffre et al., 2023)

Palaniappan et al. (2024) advocates for harmonizing healthcare AI regulations based on WHO ethical principles, while Kayusi et al. (2024) emphasizes the need for context-sensitive regulations, particularly for resource-constrained or culturally distinct populations. This highlights the challenge of balancing global standards with regional needs.

Cross-jurisdictional reviews, including those by Pesapane et al. (2021) and Palaniappan et al. (2024), stress the importance of regulatory harmonization for scalable and ethical AI deployment across borders.

Reyna et al. (2022) argue that current performance metrics fail to address real-world clinical equity, calling for a redefinition of success criteria to ensure AI benefits all populations. Additionally, "black box" systems remain largely unexplained, undermining trust in clinical decision-making, as noted by Kan (2024), Naili et al. (2025), and Chavula et al. (2024). This opacity complicates AI integration into healthcare and raises ethical concerns.

Although the EU AI Act mandates explainability, Kalodanis et al. (2025) point out that these requirements are still interpretive and untested, creating uncertainty about their real-world application and enforcement. This lack of clarity calls for more concrete frameworks to ensure AI systems are transparent, accountable, and effective in healthcare.

Policy Design

In his TEDx talk *"Want to Help Someone? Shut Up and Listen!"*, Ernesto Sirolli illustrates how well-intentioned but misaligned policy structures from Italy failed to address local needs in Zambia. Only by listening directly to local stakeholders did the mission succeed. This story holds a valuable lesson: policies must be rooted in the realities of those they aim to serve.

Drawing from the identified gaps, we observe that while policy decisions in Healthcare AI governance often reference academic studies, expert reports, and case analyses, they tend to miss two crucial components:

1. A structured, problem-solving approach that iteratively evolves from prototype to implementation to post-deployment oversight and adaptation for each medical domain like automated decision systems, chatbots, AI doctors and surgeons, AI operated medical devices.
2. Empathy-driven insights from a full spectrum of stakeholders, including national and international law enforcement agencies, WHO, the medical product/service development companies, ministries (IT and Health), medical professionals, hospital management, medical staff, researchers, engineers, advocates, institutes like National Academy of Medicine in US, Doctor's Associations, patients, and their families, and others.

Though policy committees include experts from various domains, the direct perspectives of on-the-ground stakeholders are often underrepresented. To ensure policies address real-world challenges, it is essential to foster stakeholder prioritization and manage conflicting interests between groups (e.g., medical professionals vs. AI developers). Policies must reflect local needs, particularly in resource-constrained regions, and should consider technological infrastructure disparities across jurisdictions.

Additionally, as emphasized in global regulatory reviews, it is crucial to integrate international regulatory harmonization to ensure scalable and ethical AI deployment worldwide. This will help avoid the challenges posed by fragmented regulatory approaches across different regions. Medical devices, including AI-enabled systems, are sold and utilized globally; therefore, agreements on their development, deployment, and usage should not be confined to a single jurisdiction. However, many developing countries still lack comprehensive governance frameworks for healthcare AI, resulting in regulatory disparities and potential risks in implementation.

Engaging Educational Institutes for Local Insights:

One of the most effective ways to gain region-specific insights is to engage higher educational institutions—especially those that are closely tied to local communities. Students and researchers from these institutions often have a deeper understanding of regional demographics and issues. Moreover, they can offer fresh perspectives and innovative solutions, making them valuable contributors to AI policy design.

While funding might be a challenge, many students are eager to undertake research projects to build their skills and gain hands-on experience, particularly with nationally and

internationally recognized organizations. Students from diverse academic backgrounds—such as medicine, engineering, policy, design, and psychology—can bring multi-disciplinary insights to the table. This diversity is crucial in ensuring that AI policies are holistic, inclusive, and well-rounded.

Additionally, given the rapid advancements in AI technologies, especially the rise of large language models (LLMs) like ChatGPT, it is important to ensure that policy research reflects the post-ChatGPT era. These AI technologies have widely permeated various sectors and public consciousness, making it crucial to gather new data and feedback from communities due to the shift in trust in AI. Understanding AI's widespread influence in this era will help shape policies that are relevant, timely, and forward-thinking.

Implementation Plan

Empathizing with stakeholders is a crucial first step in policy development. Conducting a comprehensive stakeholder analysis helps to understand what matters most to each group involved. This includes engaging with a diverse set of stakeholders through interviews and consented data collection. It is important to evaluate how the explainability of AI systems influences trust and adoption in real-world clinical settings. By uncovering deep, non-obvious insights that are often missed by quantitative reports, policymakers can gain a more nuanced understanding of real-world experiences. For instance, examining the impact of automated depression diagnosis on marginalized communities can provide valuable insights that guide policy decisions.

Once these insights are gathered, the next step is to clearly articulate the core problems. For example, if automated mental health diagnostics produce biased outcomes, the problem statement should reflect this issue. Success metrics need to be defined, such as reducing false negative rates, and these metrics should be adaptable to the insights gained from before. It is also important to minimize disparities across gender, caste, race, or region, ensuring that the policies are equitable and applicable to diverse populations.

To shape effective policies, it is essential to ensure that they are influenced by diverse perspectives. Before drafting the final versions, policy designers should metaphorically "wear four hats." This process includes open ideation, where solutions are generated freely; positive evaluation, where promising directions are identified; critical evaluation, where assumptions are challenged; and emotional perspective, where stakeholder feelings and ethical concerns are considered. This comprehensive approach allows for a more holistic and inclusive policy design process.

In the prototyping phase, it is crucial to build iterative policy drafts. Early versions of policies should be treated as hypotheses, not final solutions. Lightweight prototypes should be developed to test assumptions and simulate real-world application on a small scale. One valuable approach here is to introduce sandbox environments—controlled settings where domain-specific policy drafts or AI systems can be tested safely using synthetic or de-identified data. These sandboxes allow for the evaluation of outcomes, compliance, and unintended consequences without risking real patient safety or system integrity. Additionally, prototypes should include methods for interpretability and explainability to gain the trust of both patients and clinicians.

Once a prototype is ready, it is important to validate it and refine the policy drafts based on stakeholder feedback. After sandbox validation, stakeholders should be consulted again with working policy drafts for further input. It is vital to conduct Desirability, Viability, and Feasibility (DVF) assessments from major stakeholders' perspectives, such as patients, clinicians, and families. Success should be measured using both domain-specific quantitative metrics and qualitative feedback. Standardizing domain-specific fairness audits for healthcare AI tools is essential, such as requiring the submission of demographic and subgroup performance reports. Additionally, legal mandates for pre-deployment impact assessments, akin to environmental impact assessments, should be developed to ensure that the policies are thoroughly evaluated before full-scale deployment. Policies should then be refined and re-tested in cycles until they become robust.

Post-deployment oversight and adaptation are critical to ensuring that AI systems remain safe and effective over time. Mechanisms for continuous monitoring, such as periodic audits, performance reviews, our proposed SupTech solution and stakeholder surveys, should be established to track the real-world performance of the AI system. Real-world feedback loops should be integrated to identify issues like model drift, ethical violations, or fairness degradation over time. An adaptive policy framework must be created that evolves with new evidence, technological advances, or emerging social contexts. Incident reporting systems and public dashboards should be mandated for transparency and accountability in deployed AI systems. Finally, cross-sector collaboration should be encouraged to update best practices based on ongoing evaluation and international learnings.

4. Impact Assessment: Regulatory Reform for AI in Healthcare

The proposed regulatory reforms (dynamic sandbox testing, real-time performance auditing, and adaptive oversight) all represent a timely and necessary intervention to align AI innovation in healthcare with ethical and safety standards. Their implementation would significantly reshape the healthcare sector, foster more resilient global economies, and mitigate the societal risks associated with unchecked, unanticipated AI deployment.

4.1. Impact on the Healthcare Sector

For the healthcare sector, these reforms promise to close critical oversight gaps, improve patient outcomes, and enhance trust in AI-driven diagnostics and treatment tools (Bajwa et al., 2021). Dynamic sandboxes would provide a safe testing ground for emerging AI technologies, enabling clinical validation without exposing patients to undue risk (Qiu et al., 2025). This phased, real-world approach ensures that systems perform reliably across diverse populations before large-scale deployment. Real-time performance auditing would further protect patients by continuously tracking system accuracy, catching performance drifts and latent biases early. Adaptive oversight would allow for the controlled evolution of AI systems, fostering innovation while maintaining strict safety and fairness thresholds.

Collectively, these measures would establish a regulatory infrastructure that is responsive to the rapid pace of AI development. They would help clinicians integrate AI more confidently, knowing that tools are continuously evaluated and held accountable. Moreover, by addressing algorithmic bias and ensuring equitable validation, the reforms could reduce

disparities in care and improve health outcomes for historically underserved communities (Chin et al., 2025).

4.2. Impact on the U.S. and Global Economies

From an economic perspective, the proposed reforms would stimulate innovation and competitiveness in the rapidly growing AI healthtech market, projected to exceed \$200 billion globally within the next decade (Güçlü & Yıldırım, 2025). In the United States, regulatory clarity would attract responsible investment, encourage startups, and support established firms developing adaptive AI solutions. By lowering the risk of costly missteps and public trust crises, this approach would foster a more sustainable innovation ecosystem. These benefits would extend internationally as global healthtech companies seek alignment with best-in-class regulatory models (Department of Health and Social Care, 2023).

Importantly, if the U.S. leads in establishing robust, flexible standards, it could shape global norms, strengthening its position in the international AI economy. Countries lacking strong regulatory frameworks may adopt similar models, promoting cross-border trust in AI-enabled health systems (Reddy, 2024). Additionally, the improved efficiency and accuracy of AI-driven healthcare could lower system-wide costs by reducing diagnostic errors, optimizing resource allocation, and preventing unnecessary treatments, thereby easing fiscal burdens on national healthcare systems (Pavuluri et al., 2024).

4.3. Societal Impacts

On a societal level, implementing these recommendations would address the silent, systemic risks posed by healthcare AI. Transparent life cycle regulation would protect patients from harm, particularly those in marginalized communities disproportionately affected by biased or opaque algorithms (Hanna et al., 2025). By mandating fairness and traceability, the reforms would safeguard civil rights in the digital health era and promote public confidence in AI technologies used in intimate, high-stakes settings.

Equally important, these policies could help counteract digital divides between the Global North and South. By embedding equity in regulatory frameworks and potentially supporting international cooperation through WHO-aligned protocols, the reforms could encourage the development of AI tools suitable for lower-resource settings—where access to quality care is most limited.

Finally, by clearly delineating accountability among developers, providers, and institutions, these reforms would strengthen the ethical foundation of healthcare AI, ensuring that innovation enhances, not undermines, public well-being.

4.4. Conclusion

The recommended regulatory strategies offer a pragmatic yet ambitious path forward. By rebalancing innovation with oversight, these reforms will protect patients, promote equitable care, and secure the healthcare sector's transformation into the AI era, while boosting economic growth and reinforcing societal trust both in the U.S. and globally.

5. References & Appendices

References

- AlphaFold. (2025, April 2). Google DeepMind.
<https://deepmind.google/technologies/alphafold/>
- Bajwa, J., Munir, U., Nori, A., & Williams, B. (2021). Artificial intelligence in healthcare: Transforming the practice of medicine. *Future Healthcare Journal*, 8(2), e188–e194.
<https://doi.org/10.7861/fhj.2021-0095>
- Chin, M. H., Afsar-Manesh, N., Bierman, A. S., Chang, C., Colón-Rodríguez, C. J., Dullabh, P., Duran, D. G., Fair, M., Hernandez-Boussard, T., Hightower, M., Jain, A., Jordan, W. B., Konya, S., Moore, R. H., Moore, T. T., Rodriguez, R., Shaheen, G., Snyder, L. P., Srinivasan, M., Umscheid, C. A., ... Ohno-Machado, L. (2023). Guiding principles to address the impact of algorithm bias on racial and ethnic disparities in health and health care. *JAMA Network Open*, 6(12), e2345050. <https://doi.org/10.1001/jamanetworkopen.2023.45050>
- Department of Health and Social Care. (2023). Medical technology strategy. GOV.UK.
<https://www.gov.uk/government/publications/medical-technology-strategy/medical-technology-strategy>
- Fotheringham, K., & Smith, H. (2024). Accidental injustice: Healthcare AI legal responsibility must be prospectively planned prior to its adoption. *Future Healthcare Journal*, 11(3), 100181. <https://doi.org/10.1016/j.fhj.2024.100181>
- Grant, C. (2023, February 24). Algorithms are making decisions about health care, which may only worsen medical racism | ACLU. American Civil Liberties Union.
<https://www.aclu.org/news/privacy-technology/algorithms-in-health-care-may-worsen-medical-racism>
- Güçlü, M. S., & Yıldırım, E. (2025, February 10). Global market size of AI in health care to exceed \$200B in 2030. Anadolu Agency.
<https://www.aa.com.tr/en/artificial-intelligence/global-market-size-of-ai-in-health-care-to-exceed-200b-in-2030/3477366>
- Hanna, M. G., Pantanowitz, L., Jackson, B., Palmer, O., Visweswaran, S., Pantanowitz, J., Deebajah, M., & Rashidi, H. H. (2025). Ethical and bias considerations in artificial intelligence/machine learning. *Modern Pathology*, 38(3), 100686.
<https://doi.org/10.1016/j.modpat.2024.100686>
- Leckenby, E., Dawoud, D., Bouvy, J., & Jónsson, P. (2021). The Sandbox Approach and its Potential for Use in Health Technology Assessment: A Literature Review. *Applied Health Economics and Health Policy*, 19(6), 857–869. <https://doi.org/10.1007/s40258-021-00665-1>
- Medicines and Healthcare products Regulatory Agency. (2025, April 2). AI Airlock: the regulatory sandbox for AIaMD. GOV.UK.
<https://www.gov.uk/government/collections/ai-airlock-the-regulatory-sandbox-for-aiamd>

Qiu, Y., Yao, H., Ren, P., Tian, X., & You, M. (2025). Regulatory sandbox expansion: Exploring the leap from fintech to medical artificial intelligence. *Intelligent Oncology*. <https://doi.org/10.1016/j.intonc.2025.03.001>

Reader, R. (2025, January 1). The government can't ensure artificial intelligence is safe. This man says he can. *POLITICO*. <https://www.politico.com/news/2025/01/01/brian-anderson-artificial-intelligence-trump-00195782>

Reddy, S. (2024). Global harmonization of artificial intelligence-enabled software as a medical device regulation: Addressing challenges and unifying standards. *Mayo Clinic Proceedings: Digital Health*, 3(1), 100191. <https://doi.org/10.1016/j.mcpdig.2024.100191>

RFK Jr.'s health tech fix. (2025, January 30). *POLITICO*. <https://www.politico.com/newsletters/future-pulse/2025/01/30/kennedy-health-tech-fix-00201301>

Pavuluri, S., Sangal, R., Sather, J., & Taylor, R. A. (2024). Balancing act: The complex role of artificial intelligence in addressing burnout and healthcare workforce dynamics. *BMJ Health & Care Informatics*, 31(1), e101120. <https://doi.org/10.1136/bmjhci-2024-101120>

Vartan, S. (2024, February 20). Racial bias found in a major health care risk algorithm. *Scientific American*. <https://www.scientificamerican.com/article/racial-bias-found-in-a-major-health-care-risk-algorithm/>

WHO issues first global report on Artificial Intelligence (AI) in health and six guiding principles for its design and use. (2021, June 28). World Health Organization (WHO). Retrieved April 6, 2025, from <https://www.who.int/news/item/28-06-2021-who-issues-first-global-report-on-ai-in-health-and-six-guiding-principles-for-its-design-and-use>

WHO outlines considerations for regulation of artificial intelligence for health. (2023, October 19). World Health Organization (WHO). Retrieved April 6, 2025, from <https://www.who.int/news/item/19-10-2023-who-outlines-considerations-for-regulation-of-artificial-intelligence-for-health>

Xu, H., & Shuttleworth, K. M. J. (2023). Medical artificial intelligence and the black box problem: a view based on the ethical principle of “do no harm.” *Intelligent Medicine*, 4(1), 52–57. <https://doi.org/10.1016/j.imed.2023.08.001>

Kan, E. (2024). Legal architecture of the relationship between artificial intelligence and the protection of patients' rights in medical diagnostics. *Review of Law Sciences*, 8(3). <https://doi.org/10.51788/tsul.rols.2024.8.3.rjjs3425>

Naili, Y. T., Mangkunegara, I. S., Purwono, & Baballe, M. A. (2025). Regulatory challenges in AI-based diagnostics: Legal implications of AI use in medical diagnostics. *BIO Web of Conferences*, 152, 01034. <https://doi.org/10.1051/bioconf/202515201034>

Pesapane, F., Volonté, C., Codari, M., & Sardanelli, F. (2018). Artificial intelligence as a medical device in radiology: Ethical and regulatory issues in Europe and the United States. *Insights into Imaging*, 9, 745–753. <https://doi.org/10.1007/s13244-018-0645-y>

Pesapane, F., Bracchi, D. A., Mulligan, J. F., Linnikov, A., Maslennikov, O., Lanzavecchia, M. B., Tantrige, P., Stasolla, A., Biondetti, P., Giuggioli, P. F., Cassano, E., & Carrafiello, G. (2021). Legal and regulatory framework for AI solutions in healthcare in EU, US, China, and Russia: New scenarios after a pandemic. *Radiation*, 10(4), 22. <https://doi.org/10.3390/radiation1040022>

Palaniappan, K., Lin, E. Y. T., & Vogel, S. (2024). Global regulatory frameworks for the use of artificial intelligence (AI) in the healthcare services sector. *Healthcare*, 12(5), 562. <https://doi.org/10.3390/healthcare12050562>

Schulz, W., Durant, T. J. S., & Krumholz, H. (2019). Validation and regulation of clinical artificial intelligence. *Clinical Chemistry*, 65(10), 1336–1337. <https://doi.org/10.1373/clinchem.2019.308304>

Kayusi, F., Kasulla, S., Malik, S. J., Majeed, M., Turyasingura, B., Tumushabe, J. T., & Mensah, G. B. (2024). Advancing algorithmic justice: A systematic review of fair decision-making protocols in healthcare AI. *Africa Journal For Law and Development Research*, 1(1), 39–52. <https://doi.org/10.62839/ajfldr.v01.v01.39-52>

Reyna, M., Nsoesie, E., & Clifford, G. (2022). Rethinking algorithm performance metrics for artificial intelligence in diagnostic medicine. *JAMA*, 328(8), 749–750. <https://doi.org/10.1001/jama.2022.10561>

Chavula, P., Addy, A., Kayusi, F., Turyasingura, B., & Mensah, G. B. (2024). Demystifying the black box: A technical analysis of transparency requirements in clinical AI systems. *Africa Journal For Public Health Medicine & Nursing*, 1(1), 14–31. <https://doi.org/10.62839/ajfphmn.v01.v01.14-31>

Bedoya, A., Economou-Zavlanos, N. J., Goldstein, B., Young, A., Jelovsek, J., O'Brien, C., Parrish, A. B., Elengold, S., Lytle, K. S., Balu, S., Huang, E. S., Poon, E., & Pencina, M. (2022). A framework for the oversight and local deployment of safe and high-quality prediction models. *Journal of the American Medical Informatics Association*, 29(8), 1336–1342. <https://doi.org/10.1093/jamia/ocac078>

Kalodanis, K., Feretzakis, G., Anastasiou, A., Rizomiliotis, P., Anagnostopoulos, D., & Koumpourous, Y. (2025). A privacy-preserving and attack-aware AI approach for high-risk healthcare systems under the EU AI Act. *Electronics*, 14(7), 1385. <https://doi.org/10.3390/electronics14071385>

Arora, A., & Arora, A. (2022). Generative adversarial networks and synthetic patient data: Current challenges and future perspectives. *Future Healthcare Journal*, 9(2), 190–193.
<https://doi.org/10.7861/fhj.2022-0013>

Cross, J. L., Choma, M. A., & Onofrey, J. A. (2024). Bias in medical AI: Implications for clinical decision-making. *PLOS Digital Health*, 3(11), e0000651.
<https://doi.org/10.1371/journal.pdig.0000651>

Giuffrè, M., & Shung, D. L. (2023). Harnessing the power of synthetic data in healthcare: Innovation, application, and privacy. *NPJ Digital Medicine*, 6, 186.
<https://doi.org/10.1038/s41746-023-00927-3>