HARVEY CASTRO, MD, MBA
DR GPT™

AI IN EMERGENCY MEDICINE

WHAT ALL HEALTHCARE PROFESSIONALS CAN LEARN FROM THE ARTIFICIAL INTELLIGENCE REVOLUTION



WILEY

Table of Contents

- Cover
- Table of Contents
- Title Page
- Introduction
- 1 AI Is Already Here: Successes and Blind Spots in 2025
  - Executive Summary
  - Current Evidence
  - Real-World Case: Kaiser Permanente’s DAX Deployment
  - Risks and Limitations
  - Regulatory and Legal Landscape
  - Ethical Implications
  - Leadership Implementation Framework
  - Future Outlook
  - Key Takeaways
  - References
- 2 Real Patients, Real AI: From Burnout to Breakthroughs
  - Executive Summary
  - Current Evidence
  - Real-World Cases: The Kaiser Permanente Transformation
  - Risks and Limitations
  - Regulatory and Legal Landscape
  - Ethical Implications
  - Leadership Implementation Framework
  - Future Outlook
  - Key Takeaways
  - References

3 Data, Bias, and the Human Cost of Errors
- Executive Summary
- Current Evidence
- Real-World Case Studies
- Risks and Limitations
- Regulatory and Legal Landscape
- Ethical Implications
- Leadership Implementation Framework
- Future Outlook
- Key Takeaways
- References

4 Predictive Diagnosis and the End of Missed Symptoms
- Executive Summary
- Current Evidence
- Real-World Case Studies
- Risks and Limitations
- Regulatory and Legal Landscape
- Ethical Implications
- Leadership Implementation Framework
- Future Outlook
- Key Takeaways
- References

5 AI in Surgery and Imaging: Precision on a Human Timeline
- Executive Summary
- Current Evidence
- Real-World Case Studies
- Risks and Limitations
- Regulatory Fragmentation and the Cleared-but-Not-Validated Problem

Ethics, Equity, and Patient Rights
Leadership Implementation Framework
Future Outlook
Key Takeaways
References
6 Wearables and Wellness: Preventing Disease Before the ER
Executive Summary
Current Evidence
Real-World Case Studies
Risks and Limitations
Regulatory and Legal Landscape
Leadership Implementation Framework
Future Outlook
Key Takeaways
References
7 Mental Health in the Age of Algorithms
Executive Summary
Current Evidence
Real-World Cases
Risks and Limitations
Regulatory and Legal Landscape
Ethical Implications
Leadership Implementation Framework
Future Outlook
Key Takeaways
References
8 Healthcare 2030: A Vision from the Frontlines
Executive Summary
Current Evidence

Real-World Case Studies
Risks and Limitations
Regulatory and Legal Landscape
Ethical Implications for Health Equity
Leadership Implementation Framework
Future Outlook
Key Takeaways
References
9 Leading the AI Shift: What Every CEO, CIO, and CMO Must Know
Executive Summary
Current Evidence
Real-World Case Studies
Risks and Limitations
Regulatory and Legal Landscape
Ethical Implications
Leadership Implementation Framework
Future Outlook
Key Takeaways
References
10 Ethics, Equity, and AI Accountability
Executive Summary
Current Evidence
Real-World Case Studies
Risks and Limitations
Regulatory and Legal Landscape
Ethical Implications
Leadership Implementation Framework
Future Outlook
Key Takeaways

References
11 Humans at Heart, AI in Hand
Executive Summary
Current Evidence
Real-World Case Studies
Risks and Limitations
Regulatory and Legal Landscape
Ethical Implications
Leadership Implementation Framework
Future Outlook
The Closing Vision
Key Takeaways
Closing Reflection
References
Index
Copyright
Dedication
Acknowledgments
End User License Agreement

List of Illustrations

Chapter 1

Figure 1.1 Exponential growth of FDA AI medical device clearances, 1995–2025...

Figure 1.2 Hospital AI adoption of Kaiser TPMG deployment across the United St...

Figure 1.3 The AI-enabled hospital ecosystem

Chapter 2

Figure 2.1 Physician time allocation before and after ambient AI documentation

Figure 2.2 Physician burnout rate trends, 2011–2024

Figure 2.3 Ambient AI documentation restoring physician–patient connection...

Chapter 3

Figure 3.1 How biased data flows through the algorithmic pipeline to produce b...

Figure 3.2 AI diagnostic accuracy declining across the Fitzpatrick skin tone s...

Figure 3.3 Pulse oximeter racial bias: accurate readings on light skin, false ...

Chapter 4

Figure 4.1 The AI-assisted radiology workflow: scan, AI analysis, risk score, ...

Figure 4.2 Standard double-reading versus AI-assisted single-reading mammograp...

Figure 4.3 AI versus human detection performance across cancer types

Chapter 5

Figure 5.1 AI applications across six major medical specialties

Figure 5.2 The oncology AI continuum: screening through survivorship

Chapter 6

Figure 6.1 Manual workflow versus AI-simplified workflow

Figure 6.2 AI-automated billing flow: encounter to payment

Figure 6.3 Institutional ROI from AI implementation: time and cost recovery

Chapter 7

Figure 7.1 The mental health treatment gap: those with illness versus those re...

Figure 7.2 AI mental health tool safety spectrum: low-risk wellness to high-ri...

Figure 7.3 Therabot RCT outcomes: before and after reductions in depression, a...

Chapter 8

Figure 8.1 The physician shortage versus AI capacity expansion

Figure 8.2 AI capability roadmap 2025–2030: reactive AI to predictive AI

Figure 8.3 Reactive AI (linear) versus agentic AI (multidirectional)

Chapter 9

Figure 9.1 AI adoption S-curve: pilot to early adoption to majority to optimiz...

Figure 9.2 AI governance committee organizational structure

Figure 9.3 AI vendor evaluation scorecard: evidence, bias, governance, oversig...

Chapter 10

Figure 10.1 The AI adoption equity gap across the United States: rural versus ...

Figure 10.2 Four pillars of ethical AI in healthcare: (1) fairness/bias mitiga...

Figure 10.3 The AI liability chain: unanswered accountability questions at eve...

Chapter 11

Figure 11.1 Human cognition and AI cognition: the hybrid advantage zone

Figure 11.2 The gift of time: from documentation burden to patient presence

Figure 11.3 The full arc: from reactive AI to human-AI partnership

AI in Emergency Medicine

What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution

Harvey Castro, MD, MBA DR GPT™

WILEY

Introduction

I have a scar on my left hand. It is faded now, but I can still trace the ridge where the tissue healed without care, stitched by a doctor who never looked me in the eye, on a night when a 12-year-old boy and his Colombian immigrant mother waited eight hours in an emergency room and were processed, not treated.

That experience is where this book begins. Not because it's dramatic, though it was, but because it's not exceptional. Across the United States, millions of patients every year feel what I felt that night: invisible. Overwhelmed by a system that was never designed for them, triaged by workflows that move too fast to see individuals, and failed by the gap between what medicine knows how to do and what it actually delivers at the point of care.

I became an emergency physician specifically to close that gap. I ran ERs. I built more than 30 healthcare applications: apps in 54 languages, coloring book discharge instructions for children, sign language video guides for patients who are hard of hearing. I taught myself to code between patients on night shifts because a nurse once flipped through a physical textbook while we were coding a patient and something in my mind refused to accept that this was the best we could do.

AI entered my life not as a fascination but as a necessity. I needed tools that didn't get tired at 3 a.m. Tools that didn't miss drug interactions because the chart was buried in 1,000 pages of EHR notes. Tools that wouldn't let a 12-year-old boy bleed for eight hours because the system was overwhelmed and no one was watching.

This book is the result of everything I have learned at the intersection of emergency medicine, healthcare technology, and artificial intelligence. It is grounded in evidence: current through early 2026 and drawn from landmark clinical trials, peer-reviewed literature, and institutional case studies from the institutions building this future right now. It is also grounded in something the data alone can't convey: the understanding that every statistic in these pages represents a patient, and every patient deserves to be seen.

My mother survived two cancers because she was an ultrasound technician who happened to scan herself. My father died of a sudden cardiac event with no warning signs and no early detection. One parent had access to the technology. The other didn't. That is not a healthcare system. It is a lottery. And it is the lottery I have spent my career trying to rig in the patient's favor.

That is the urgency behind AI in Emergency Medicine. Not the market projections, though they are extraordinary. Not the regulatory milestones, though they matter. The urgency is human. AI is already saving lives in the hospitals that have deployed it well. It is already failing patients in the systems that have not asked the hard questions about bias, accountability, and equity. The difference between those two outcomes is not the algorithm. It is the leadership, the governance, and the willingness to hold the technology to the same standard we hold the physician: first, do no harm.

Human plus AI beats the best human or the best AI alone. The evidence is clear on this. The human has to be in the room, in the decision, and in the values that guide every deployment. That is the rung I am reaching for. I hope this book helps you reach for it too.

Harvey Castro, MD, MBA DR GPT™
Dallas, Texas
March 2026

1
AI Is Already Here: Successes and Blind Spots in 2025

I was 12 years old when I learned what it feels like to be invisible inside a hospital.

A rusty, jagged signpost did the damage, slicing my hand open on a summer afternoon in a neighborhood where signposts were never maintained and kids like me played in the street because there was nowhere else to go. My mother grabbed a towel, wrapped my hand, and drove me to the nearest emergency room. Blood soaked through the towel before we parked.

We waited eight hours. Eight hours of holding a blood-soaked cloth, watching the clock on the wall tick past dinner, past bedtime, past any reasonable definition of care. When the doctor finally appeared, he stitched me up without pain medication. No lidocaine. No numbing agent. Nothing. Just needle, thread, and a 12-year-old boy gripping the exam table with his free hand, trying not to scream.

The pain was terrible. But what lingered far longer was the feeling that I was just a number. Another chart in a stack. Another body on a gurney. The doctor never looked me in the eye. My mother, an immigrant from Colombia who worked multiple jobs and stretched food stamps to feed us both, didn’t have the words to demand better. We were processed. We were not cared for.

That scar is still on my hand. I look at it every day. For years, it pushed me away from medicine. Why would I enter a system that treated people like that? But eventually, the scar pulled me back. I became an ER doctor specifically so that no one else would feel invisible the way my mother and I did.

Decades later, standing in my own emergency department, I watched a monitor do something that changed the trajectory of my thinking. The hospital had deployed a predictive sepsis algorithm; it was nothing glamorous, just a model running quietly behind the vitals screen. A patient

had come in with vague complaints: fatigue, mild fever, a general sense of something wrong. Vitals looked borderline normal. By every traditional metric, this patient could wait.

The algorithm disagreed. It flagged the case: a subtle trend in heart rate variability and a lactate trajectory the human eye would miss in a busy shift. I ordered labs. The lactate came back elevated. We caught the sepsis hours before the classic cascade would have made it obvious. In sepsis, these hours are the difference between a patient walking out and a patient leaving in a body bag.

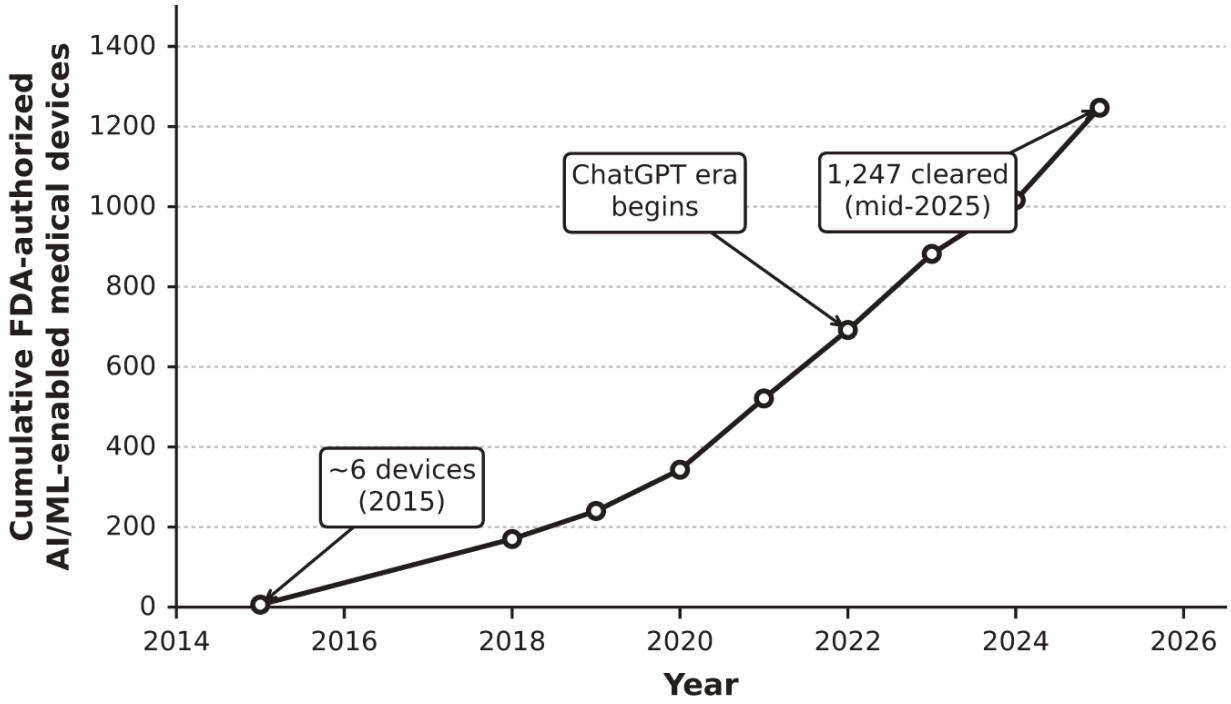
I stood in the hallway after that case and stared at the scar on my hand. Twelve-year-old me waited eight hours because the system didn't see him. This patient didn't wait because a machine saw what a busy human might have missed. That was the moment I understood: artificial intelligence (AI) wasn't going to replace the physician's touch. It was going to give us the time and the clarity to provide it. It is that touch that was missing the night a boy bled through a towel in a waiting room and no one came.

Executive Summary

Artificial intelligence is no longer a future promise in healthcare; it is an operational reality reshaping how we diagnose disease, document encounters, predict deterioration, and deliver care. By mid-2025, the U.S. Food and Drug Administration had authorized more than 1,250 AI-enabled medical devices, a nearly fifty-fold increase from just six devices a decade earlier (see Figure 1.1) [1]. Two out of three American physicians reported using AI in clinical practice, up from barely one in three the year before [2]. The global healthcare AI market had swelled to nearly $15 billion, with projections pointing toward $110 billion by the decade's end [3]. These are not incremental shifts. They are the opening chapters of a structural transformation.

Yet the story of AI in medicine is not a simple triumph narrative. For every landmark success—mammography AI detecting 29 percent more cancers in the first randomized controlled trial of its kind [4], ambient documentation tools pulling thousands of physicians back from the brink of burnout [5], sepsis algorithms catching deterioration twelve hours before human

clinicians [6]—there is a cautionary counterpoint. IBM poured billions into Watson Health and walked away with little to show for it [7]. A widely deployed sepsis model in one of America's largest electronic health record systems missed two-thirds of sepsis cases in independent validation [8]. Racial bias embedded in a risk-prediction algorithm affected how care was allocated to 200 million Americans [9]. And for the first time in its history, ECRI, the nation's foremost health technology safety authority, named AI the number-one health technology hazard heading into 2025 [10].



Source: FDA AI/ML-Enabled Medical Devices list, accessed mid-2025.

Figure 1.1 Exponential growth of FDA AI medical device clearances, 1995–2025

This chapter maps the territory as it exists: the genuine victories, the persistent blind spots, the regulatory frameworks struggling to keep pace, and the ethical imperatives that will define whether AI becomes medicine's greatest ally or its most consequential liability. As an emergency physician who has watched these technologies arrive at the bedside, I can tell you that the revolution is not coming. It is here. The question now is whether we are ready to govern it wisely.

Current Evidence

The acceleration of AI in healthcare is best understood through three lenses: regulatory clearances, market investment, and clinical adoption. Each tells a complementary story of a technology ecosystem reaching critical mass.

FDA Clearances: From Trickle to Torrent

In 2015, the Food and Drug Administration (FDA) authorized six AI-enabled medical devices. By the end of 2025, that number had climbed past 1,250, a trajectory more like a hockey stick than a gentle growth curve (flat progress for years and then a sudden, steep climb) [1]. The 2024 cohort alone included 168 devices, with radiology accounting for 74.4 percent and cardiovascular applications comprising 6.5 percent [11]. The 2025 cohort surged further to 295 clearances from 221 unique manufacturers, with a median clearance time of 142 days [12]. In 2024, the overwhelming majority (94.6 percent) came through the 510(k) pathway, the FDA's most common route for devices that demonstrate substantial equivalence to existing products [11].

A telling detail emerged from the 2024 data: for the first time, non-U.S. applicants accounted for the majority of approvals at 57.7 percent, with France, China, South Korea, and Israel joining the United States as top contributors [11]. This globalization of medical AI development means the competitive landscape is no longer a domestic conversation. It is an international race.

The FDA's December 2024 finalization of its Predetermined Change Control Plan guidance marked a pivotal regulatory milestone, establishing a formal framework for AI devices to evolve through prespecified updates without requiring entirely new marketing submissions [13]. This was a recognition that AI is fundamentally different from traditional medical devices, that it improves over time, and that regulatory frameworks must accommodate that reality.

Market Investment: Following the Money

The financial trajectory mirrors the regulatory one. The global healthcare AI market reached approximately $14.92 billion in 2024, with

MarketsandMarkets research firm projecting growth to $110.61 billion by 2030 [3]. In the startup ecosystem, the signal was even more dramatic: Rock Health reported that AI-focused companies captured 54 percent of all digital health venture funding in 2025, a striking consolidation of investor attention around a single technology paradigm [14].

These numbers carry weight because they represent institutional bets, not speculative enthusiasm. Still, capital is being deployed by venture firms, health systems, and pharmaceutical companies that are staking organizational futures on AI's trajectory.

Clinical Adoption: From Early Adopters to Mainstream

Perhaps the most revealing metric is the pace of clinician adoption. The American Medical Association's 2024 survey found that 66 percent of physicians reported using AI in their clinical practice, a dramatic increase from 38 percent in 2023 [2]. On the institutional side, the Office of the National Coordinator for Health Information Technology reported that 71 percent of nonfederal acute-care hospitals were using predictive AI as of 2024 [15]. This is no longer early-adopter territory. It is mainstream.

But adoption numbers alone can be misled. The same AMA survey revealed that 47 percent of physicians wanted more oversight of AI tools [2], and a 2024 study in JMIRx Med found that 72.4 percent of healthcare workers expressed concern about AI-driven de-skilling, the fear that reliance on algorithms would erode the clinical judgment developed over years of training [16]. Higher AI literacy, paradoxically, was associated with greater skepticism, not less [17]. The physicians who understood the technology best were the ones most worried about its limitations.

Where AI Is Winning: Landmark Successes

Several clinical applications have moved beyond pilot programs into routine care, backed by large trials and real-world deployment data. The following cases represent the strongest evidence to date.

AI in Radiology: The MASAI Trial and Beyond If there is a single study that marked 2025 as the year AI in medicine came of age, it may be the MASAI trial. *Published in The Lancet in January 2026,* MASAI was

the first completed randomized controlled trial of AI-supported mammography screening, and its results were extraordinary [4].

The trial enrolled 105,915 women across Swedish screening sites and randomized them to either standard double reading by two radiologists or AI-supported reading, with AI serving as the first reader. The results: AI-supported screening detected 29 percent more cancers than standard practice, reduced interval cancers (the ones that appear between scheduled screenings) by 12 percent, and achieved a 44 percent reduction in radiologist screen-reading workload. Critically, there was no increase in false-positive rates [4]. This was not a marginal improvement. It was a paradigm shift, demonstrating that AI could simultaneously improve diagnostic accuracy and reduce the human labor required to achieve it.

The MASAI results landed in a radiology landscape already transformed by AI. By 2024, 74.4 percent of all FDA-cleared AI devices were in radiology [11], and the specialty had become the proving ground for AI's clinical value proposition.

In stroke care, the evidence base was equally compelling. Brainomix 360, an AI-powered stroke imaging platform deployed across 107 NHS hospitals in the United Kingdom, demonstrated a 100 percent increase in endovascular thrombectomy procedures (a clot-removal surgery that must be performed within hours of a stroke; more procedures means more patients treated in time) and a 64-minute reduction in door-in-door-out time across a dataset of more than 450,000 patients, as reported in *The Lancet Digital Health* in December 2025 [18]. The clinical significance of 64 fewer minutes in stroke care is difficult to overstate. In acute ischemic stroke, approximately 1.9 million neurons are lost per minute of delayed treatment.

Meanwhile, Viz.ai, a platform that provides AI-powered stroke triage, underwent a rigorous meta-analysis. Published in *Neurology* in April 2025, the analysis of 15,402 patients demonstrated significant reductions in time-to-treatment across multiple stages of the stroke care pathway [19]. These were not theoretical improvements demonstrated in controlled settings. They were real-world outcomes on a scale.

I have seen this transformation firsthand. In the emergency department (ED), seconds matter in stroke care. The difference between a patient receiving endovascular therapy within the treatment window and missing

that window can be the difference between walking out of the hospital and spending months rehabilitating. When I receive an AI-generated alert on my phone that a large vessel occlusion has been detected before the patient even arrives in my ED, the entire care trajectory shifts. Radiology teams are mobilized. Interventionalists are notified. The patient enters a system already primed to respond. That is not incremental improvement; it is a fundamentally different model of emergency care.

Ambient AI Documentation: The Burnout Revolution If AI in radiology is the headline success of diagnostic medicine, ambient AI documentation is the story that resonates most deeply with frontline clinicians. The reason is simple: physician burnout has reached crisis proportions, and documentation burden is its primary driver.

The numbers tell a devastating story. Before ambient AI, physicians spent approximately two hours on documentation for every one hour of direct patient care. The electronic health record, intended to improve care coordination, had instead become a documentation trap. This system required clinicians to spend their evenings and weekends completing notes rather than resting, recovering, or engaging with their families.

Ambient AI changed the equation. These systems, powered by products such as Nuance's Dragon, use large language models to listen to patient-clinician conversations and automatically generate clinical documentation. The physician reviews and signs the note. The hours of after-hours typing vanish.

The evidence base for these tools matured dramatically in 2024 and 2025. A randomized controlled trial published in NEJM AI in November 2025 enrolled 238 physicians at UCLA and found that both AI documentation tools tested significantly improved burnout scores [5]. A separate RCT at Providence Health, published in the *Future Healthcare Journal* in July 2025, reported a 30.3 percent decrease in burnout and a 2.5 hours-per-week reduction in documentation time [20]. A multicenter quality improvement study published in JAMA documented a decrease in physician burnout from 51.9 percent to 38.8 percent within just 30 days of ambient AI deployment [21].

Dragon Copilot alone had surpassed 600,000 clinician users across more than 400 healthcare organizations by 2025 [22], making it one of the most widely deployed AI tools in the history of medicine. The scale was enormous, and the signal was consistent: ambient documentation AI reduced burnout, saved time, and improved the clinician experience.

As an emergency physician, I can tell you that the documentation burden in my specialty is relentless. Each patient encounter requires a history and physical exam, medical decision-making documentation, procedure notes, and discharge instructions, often while simultaneously managing other critically ill patients. When I first used an ambient AI tool, I finished my shift with my notes complete. I did not take work home. I did not log into the electronic health record (EHR) at midnight. That experience, multiplied across hundreds of thousands of clinicians, represents a genuine quality-of-life revolution in medicine.

Sepsis Prediction: COMPOSER, TREWS, and the Race Against Time Sepsis kills approximately 350,000 Americans annually and remains the leading cause of in-hospital mortality. Every hour of delayed treatment increases the odds of death by approximately 4 to 8 percent. The clinical challenge is that sepsis often presents with subtle, nonspecific signs that are easy to miss in a busy hospital until the patient is crashing.

AI sepsis prediction systems represent one of the most consequential applications of machine learning in acute-care medicine. The COMPOSER system, evaluated in a study published in *npj Digital Medicine* in January 2024, demonstrated a 17 percent reduction in sepsis-related mortality and 82 percent sensitivity for early detection [23]. TREWS, developed at Johns Hopkins University, achieved an area under the curve (AUC) of 0.94, where AUC measures how well a model distinguishes between two outcomes on a scale of 0.5 (no better than chance) to 1.0 (perfect accuracy), indicating near-perfect discrimination, and could detect sepsis up to 12 hours before clinical recognition, with 82 percent sensitivity [6].

In January 2026, the HERACLES study, spanning 97,559 hospital stays, demonstrated significant reductions in mortality associated with AI-driven sepsis detection, adding further weight to the evidence base [24]. These studies collectively suggest that well-validated AI sepsis tools can save

lives by reducing the time from the first physiological signal to clinical intervention.

The sepsis story is particularly important because it illustrates both the promise and the peril of medical AI. As we will see later in this chapter, not all sepsis algorithms perform as advertised. The gap between what is possible with rigorously validated systems and what is deployed in practice is one of the most critical tensions in the field.

Diagnostic AI Across Specialties Beyond radiology and sepsis, AI diagnostic tools demonstrated value across a widening range of clinical specialties in 2024 and 2025. In dermatology, AI-powered skin lesion analysis tools achieved diagnostic accuracy comparable to that of board-certified dermatologists for melanoma detection. In ophthalmology, autonomous AI systems for diabetic retinopathy screening initially cleared by the FDA in 2018 continued to expand into community health settings, bringing specialist-level screening to populations that had historically lacked access. In pathology, AI-assisted whole-slide imaging analysis shifted from research novelty to clinical workflow integration, with multiple FDA-cleared tools entering routine use for prostate cancer grading and breast cancer biomarker assessment.

Cardiovascular space saw notable growth, rising from 6.5 percent of FDA clearances in 2024 to 8.8 percent in 2025 [11,12], with AI-powered ECG interpretation, cardiac MRI analysis, and atrial fibrillation detection tools all reaching clinical deployment. The common thread across these specialties was pattern recognition AI excelling at tasks that required consistent analysis of large volumes of visual or signal data, precisely the tasks where human performance degrades with fatigue and volume.

The Blind Spots: What the Headlines Miss

For every success story, there is a cautionary tale. The failures in healthcare AI have been as instructive as the victories, often more so. Three cases stand out for what they reveal about the gap between promise and practice.

IBM Watson: The $4 Billion Cautionary Tale No discussion of AI in healthcare is complete without reckoning with Watson Health, perhaps the most expensive failure in the history of health technology. IBM invested an

estimated $4 to $5 billion in its Watson Health division, acquiring Truven Health Analytics, Phytel, Explorys, and Merge Healthcare in a rapid acquisition spree that promised to bring cognitive computing to cancer treatment, drug discovery, and population health management [7].

The vision was seductive: Watson would ingest the world's medical literature, process patient records, and deliver personalized treatment recommendations that rivaled the judgment of the world's best oncologists. IBM's marketing was breathtaking in its ambition. The reality was humbling.

Internal IBM documents later revealed by STAT News showed that Watson for Oncology frequently generated treatment recommendations that, in the words of IBM's own engineers, were "unsafe and incorrect" [7]. The system struggled with the complexity of real-world clinical decision-making, including nuance, comorbidities, patient preferences, and treatment contraindications, which experienced oncologists navigate through years of training and pattern recognition. Watson's recommendations were often based on the preferences of a small number of Memorial Sloan Kettering physicians whose expertise had been used to train the system, rather than on the broad evidence base the technology had promised to synthesize.

By 2022, IBM had sold the Watson Health data and analytics assets to Francisco Partners for an estimated $1 billion, a fraction of the original investment [7]. The Watson story is not merely a corporate failure. It is a warning about what happens when AI hype outpaces technical capability, when marketing promises are substituted for clinical validation, and when the inherent complexity of medicine is underestimated.

The lesson for today's AI implementers is clear: clinical AI must be validated in the specific populations and workflows where it will be deployed, not simply demonstrated in controlled research environments. The gap between a research prototype and a reliable clinical tool is where billions of dollars and potentially patient lives can be lost.

Epic Sepsis Model: When Proprietary Goes Wrong If Watson was a failure of ambition, the Epic Sepsis Model was a failure of transparency. The Epic Sepsis Model was integrated into Epic Systems' electronic health record platform—the most widely used EHR in the United States—and

deployed across hundreds of hospitals. Epic reported the model's area AUC of 0.76 to 0.83 (a range indicating moderate to good predictive accuracy), suggesting the ability to reliably distinguish patients who would deteriorate from those who would not [8].

In 2021, a landmark study published in *JAMA Internal Medicine* by researchers at the University of Michigan told a very different story. Their independent external validation found the model's actual AUC was 0.63, barely better than chance, and that it missed 67 percent of sepsis cases [8]. The model generated so many false alerts that clinicians developed "alert fatigue" and routinely dismissed the warnings. A tool designed to save lives had become background noise.

The Epic Sepsis Model controversy exposed several systemic vulnerabilities. First, the model was developed and validated using Epic's internal datasets, with limited independent external validation before widespread deployment. Second, the proprietary nature of the algorithm meant that hospitals deploying it could not fully interrogate its methodology, training data, or performance characteristics. Third, the gap between the developer's reported performance metrics and real-world independent validation highlighted the critical importance of external validation. This step remains disturbingly optional in the current regulatory landscape.

The broader lesson is institutional: when health systems deploy proprietary AI tools embedded in their EHR platforms, they are often trusting the vendor's performance claims without the ability to verify them. This is a governance failure, not merely a technical one.

The Evidence-Adoption Gap: What the JAMA Summit Found In October 2025, the JAMA Health Policy Summit convened experts to assess the state of AI implementation in healthcare. Their conclusion was sobering: most AI tools enter practice with minimal testing [25]. Despite the proliferation of FDA-cleared devices, the evidence base for many deployed tools remained thin. Clearance, the experts noted, is not the same as validation. A 510(k) clearance demonstrates substantial equivalence to a predicate device (a previously cleared product that a new device claims to be substantially equivalent to for regulatory approval)—it does not require a randomized controlled trial demonstrating clinical benefit.

Data on device safety further illuminated this evidence-adoption gap. A 2025 study in *JAMA Health Forum* found 182 AI device recall events, with 43 percent occurring within 1 year of market entry [26]. These recalls, triggered by software malfunctions, incorrect outputs, and performance degradation, underscore that market clearance does not guarantee ongoing reliability.

The paradox is striking in an era when more AI tools are available than ever; the infrastructure for evaluating their real-world performance remains inadequate. Post-market surveillance for AI devices is fragmented, registries for tracking outcomes are nascent, and the burden of monitoring falls disproportionately on health systems that often lack the technical expertise to monitor effectively.

Real-World Case: Kaiser Permanente's DAX Deployment

The most comprehensive real-world deployment of ambient AI documentation in the United States provides a detailed case study in both the potential and the complexity of scaling AI across a large health system.

The Permanente Medical Group (TPMG), the physician group of Kaiser Permanente in Northern California, deployed an ambient AI documentation tool (initially Nabla, and later transitioned to Abridge as part of Kaiser Permanente's enterprise-wide rollout in 2024) across its organization, which became one of the largest clinically integrated AI implementations in the country. The results, published in *NEJM Catalyst* in 2025, are remarkable in their scale and detail [27].

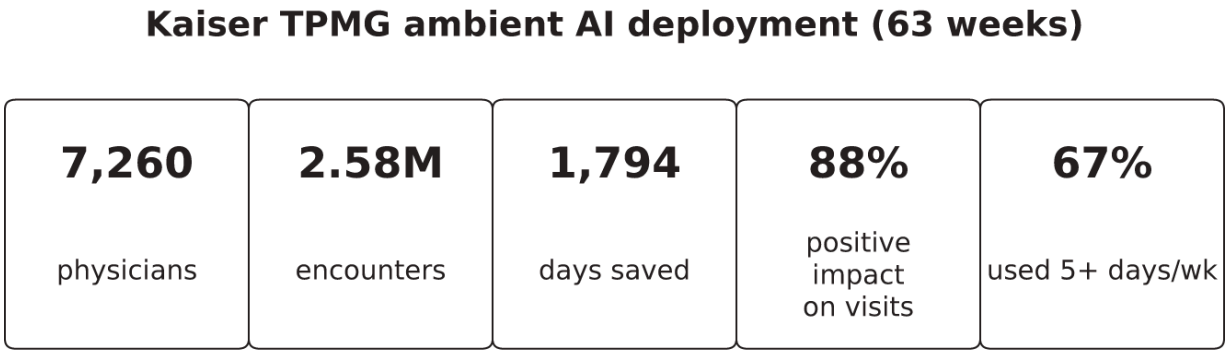
- 7,260 physicians actively using the system
- More than 2.5 million patient encounters processed through ambient AI documentation
- 1,794 working days saved in cumulative physician documentation time, the equivalent of nearly five physician-years recovered from documentation burden
- 84 percent physician satisfaction with the technology

The TPMG deployment offers several lessons for health systems considering adopting ambient AI. First, scale matters (see Figure 1.2). TPMG did not run a pilot with 50 physicians and declare success. They rolled out across more than 7,000 clinicians and measured outcomes across millions of encounters. Second, the time savings were not theoretical—they were measured in actual working days recovered, which was time that physicians could redirect to direct patient care, education, or personal restoration. Third, the satisfaction scores suggest that this technology addressed a genuine pain point rather than creating a new one.

The clinical implications extend beyond individual physician well-being. Documentation quality improved as AI-generated notes captured details of the conversation that hurried physicians might otherwise omit. Patient-physician interactions improved as clinicians maintained eye contact and engaged in conversation rather than typing on a computer screen during the visit. And the organizational economics were favorable: recovered physician time translates directly into either increased patient access or reduced overtime costs.

The TPMG experience also illustrates the challenges of implementation. Successful deployment required significant change management, training, and workflow redesign. Not all specialties were adopted at the same pace. Physicians needed to develop new skills learning to review and edit AI-generated documentation rather than creating it from scratch. Ongoing monitoring was required to ensure documentation accuracy and regulatory compliance.

For health system leaders evaluating ambient AI, the TPMG case provides the most robust evidence to date that these tools deliver measurable value at scale. But it also demonstrates that success requires organizational commitment, not just technology procurement.



Source: Tierney et al., NEJM Catalyst (Mar 2025).

Figure 1.2 Hospital AI adoption of Kaiser TPMG deployment across the United States

Risks and Limitations

While AI in emergency medicine holds great promise, it's important to also consider the risks and limitations.

Algorithmic Bias: The Deepest Vulnerability

The most consequential risk of AI in healthcare is also the most insidious: algorithmic bias that systematically disadvantages the patients who are already most vulnerable.

The landmark 2019 study by Obermeyer and colleagues, published in *Science*, exposed how a widely used commercial algorithm had baked racial bias into its very architecture [9]. The algorithm used healthcare costs as a proxy for healthcare needs, a seemingly reasonable design choice that carried a devastating flaw. Because Black patients in the United States historically receive fewer healthcare dollars than white patients with equivalent illness severity, the algorithm systematically scored Black patients as healthier than they actually were. At any given risk score, Black patients had 26.3 percent more chronic conditions than white patients. The algorithm affected care allocation decisions for approximately 200 million people annually [9].

The pulse oximetry problem revealed bias in an even more fundamental medical device. A 2020 study by Sjoding and colleagues, published in the

New England Journal of Medicine, found that Black patients experienced hidden hypoxemia oxygen levels that appeared normal on pulse oximetry but were actually dangerously low at approximately three times the rate of white patients [28]. This was not an AI error in the traditional sense. Still, it illustrates how training data and device calibration can embed racial disparities into the tools clinicians rely on every day.

In 2025, the bias problem extended to large language models. A systematic review found racial bias present in 90.9 percent of evaluated studies and gender bias in 93.7 percent [29]. A Cedars-Sinai study published in *npj Digital Medicine* found that leading LLMs made dramatically different psychiatric treatment recommendations based solely on a patient's stated or implied race in one case, omitting medication recommendations for attention-deficit/hyperactivity disorder (ADHD) when the patient was identified as Black but including them when race was not specified [30].

Perhaps most alarming was the transparency deficit: a 2025 systematic review published in *Oxford Open Digital Health* found that 84 percent of clinical AI studies globally did not report the racial or ethnic composition of their training data [31]. We cannot fix bias we cannot see, and we cannot see bias if we do not report who is and who is not represented in the data that trains these systems.

The Black Box Problem

Many AI systems, particularly deep learning models, operate as functional black boxes, producing outputs without interpretable explanations of how those outputs were generated. For a physician deciding whether to act on an AI recommendation, this opacity creates a fundamental trust challenge. A radiologist can evaluate a colleague's reading by reviewing the same images and reasoning. A clinician cannot do the same with a neural network that has processed millions of parameters through dozens of hidden layers.

The explainability problem is not merely philosophical. It has direct clinical consequences. When a sepsis alert fires and the physician cannot understand why and the algorithm's reasoning cannot be interrogated, the alert becomes an interruption rather than a decision-support tool. Alert fatigue follows. Life-saving warnings get dismissed because clinicians cannot distinguish true positives from noise.

Regulatory bodies have recognized this gap. The FDA's 2024 transparency guiding principles for ML-enabled devices represented an initial step, and the EU AI Act's requirements for high-risk AI systems include explainability provisions. But the technical reality is that the most powerful AI models are often the least interpretable, creating a tension between performance and transparency that the field has not yet resolved.

Data Quality and EHR Gaps

AI is only as good as the data it is trained on, and healthcare data is notoriously messy. Electronic health records are designed for billing and documentation, not for machine learning. Diagnoses are frequently coded for reimbursement purposes rather than clinical accuracy. Free-text notes contain abbreviations, misspellings, and idiosyncratic shorthand that vary by institution, specialty, and individual physician. Vital signs may be entered hours after they were measured. Lab results may reflect ordered tests rather than clinical suspicion.

These data quality issues cascade through the AI development pipeline. A model trained on biased billing codes will learn biased associations. A model trained on data from a single academic medical center may fail when deployed in a community hospital with a different patient population, documentation culture, and workflow. The Epic Sepsis Model's performance degradation in external validation was, at least in part, due to data quality and generalizability issues [8].

The fragmentation of U.S. healthcare data compounds the challenge. Patients receive care across multiple health systems, each with its own EHR, data standards, and interoperability limitations. The AI that sees only one slice of a patient's history faces incomplete information, a challenge that even the most sophisticated algorithms cannot overcome solely through architecture.

The Clinician Trust Deficit

Trust is the currency of clinical AI adoption, and the current account balance is precarious. The AMA's 2024 survey revealed a striking duality: 66 percent of physicians were using AI, but 47 percent wanted more oversight of these tools [2]. This is not resistance to technology. It is

informed caution from professionals who understand that the stakes in medicine are measured in human lives.

A 2024 study in *JMIRx Med* found that 72.4 percent of healthcare workers expressed concern about AI-driven de-skilling and the erosion of clinical judgment that might result from excessive reliance on algorithmic recommendations [16]. This fear is not irrational. Skills atrophy without practice, and if a generation of clinicians learns to defer to algorithms rather than developing independent diagnostic reasoning, the consequences of system failures become catastrophic.

The trust paradox identified in a 2025 study published in *BMC Medical Ethics* added another dimension: physicians with higher AI literacy were actually more skeptical of AI tools, not less [17]. Knowledge bred caution. The clinicians who understood the limitations of machine learning sensitivity to training data, the potential for distribution shift, and the absence of true causal reasoning were precisely the ones most wary of deploying these tools without adequate safeguards.

The Digital Divide

The promise that AI will democratize healthcare access confronts an uncomfortable reality: the communities that stand to benefit most from AI-augmented care are often the least likely to receive it. A January 2026 analysis published in Nature Health found that rural and underserved hospitals were significantly less likely to adopt AI tools [32]. The reasons are structurally limited broadband infrastructure, smaller IT budgets, fewer data scientists on staff, and lower volumes of training data.

The Federal Communications Commission estimates that 35 percent of rural Americans lack adequate broadband connectivity [33], a prerequisite for cloud-based AI tools that require real-time data transfer and processing. This creates what researchers have termed a “performance paradox”: AI systems work worst for the populations that need them most, because those populations are underrepresented in training data, underserved by digital infrastructure, and under-resourced for implementation [32].

If health systems in well-funded urban academic centers deploy AI tools that improve outcomes for their patient populations while rural and safety-

net hospitals are left behind, AI will not have narrowed health disparities. It will have widened them.

Regulatory and Legal Landscape

The regulatory framework for healthcare AI is evolving rapidly, but it remains a patchwork of enacted legislation, finalized guidance, proposed rules, and conspicuous gaps. Leaders must understand not only what regulations exist but where the most consequential gaps remain.

FDA Framework

The FDA framework includes the following:

> **ENACTED/FINAL GUIDANCE:** The FDA has authorized more than 1,250 AI-enabled medical devices through mid-2025 [1], primarily through the 510(k) pathway. The agency's most significant recent action was the finalization of its Predetermined Change Control Plan (PCCP) guidance on December 4, 2024 [ENACTED] [13]. This guidance, rooted in the Food and Drug Omnibus Reform Act of 2022, allows manufacturers to prespecify planned algorithm modifications in their initial marketing submissions, enabling iterative AI improvement without requiring a new clearance for each update. By late 2025, approximately 10 percent of new AI device clearances included PCCPs [12].
>
> The FDA also continued to implement its AI/ML-Based Software as a Medical Device Action Plan [ENACTED 2021, ongoing implementation], which established a five-pronged strategy covering regulatory frameworks, good machine learning practices, patient-centered approaches, algorithm transparency, and real-world performance monitoring.
>
> **Key limitation:** As of mid-2025, no FDA-authorized AI medical device used generative AI or large language models [1]. The regulatory framework for LLM-based clinical tools remains undefined, creating uncertainty for developers of ambient documentation, clinical decision support, and diagnostic tools that use foundation models.

EU AI Act

The EU's AI Act specifies the following:

> **ENACTED:** As of August 1, 2024, the European Union's Artificial Intelligence Act represents the world's most comprehensive AI regulatory framework [34]. Under the act, medical AI systems are classified as high-risk, subjecting them to mandatory requirements including:
>
> - Conformity assessments before market entry
> - Risk management systems throughout the AI lifecycle
> - Data governance and quality requirements for training datasets
> - Technical documentation and transparency obligations
> - Human oversight requirements
> - Post-market monitoring obligations
>
> Implementation follows a phased timeline: prohibited AI practices took effect in February 2025, high-risk system requirements begin phasing in through 2026 and 2027, and full enforcement across all provisions is expected by 2027. The act imposes fines of up to 35 million euros or 7 percent of global annual turnover for violations [34].
>
> For U.S.-based AI developers serving European markets, the EU AI Act creates extraterritorial compliance obligations that will reshape product development, documentation practices, and data governance.

WHO Guidance

WHO has published the following guidance:

> **PUBLISHED:** In January 2024, the World Health Organization (WHO) issued guidance on the ethics and governance of large multimodal models in healthcare, addressing foundation models (including GPT-4, Med-PaLM, and similar systems) [35]. The guidance outlined five key areas requiring attention: transparency, accountability, inclusiveness, safety, and data privacy. While not legally binding, WHO guidance carries significant influence in shaping

national regulatory approaches, particularly in low- and middle-income countries.

Joint Commission/CHAI Guidance

The Joint Commission published the following:

> **PUBLISHED:** In September 2025, the Joint Commission, in collaboration with the Coalition for Health AI (CHAI), published guidance on AI governance for accredited healthcare organizations [36]. This guidance established expectations for health system AI governance committees, risk assessment frameworks, clinical validation requirements, and ongoing monitoring obligations. For the thousands of U.S. hospitals that depend on Joint Commission accreditation, this guidance effectively created a de facto governance standard even in the absence of federal legislation.

Key Regulatory Gaps

Despite these advances, critical gaps remain:

> **No federal algorithmic audit mandate:** There is no U.S. federal requirement for independent auditing of clinical AI algorithms, either before or after deployment. Manufacturers self-report performance metrics, and independent validation remains voluntary.
>
> **AI liability framework unclear:** When an AI tool contributes to a misdiagnosis or adverse outcome, the legal framework for liability allocation among the AI developer, the health system deploying the tool, and the clinician acting on the recommendation remains unsettled. No federal legislation has established clear liability standards for AI-assisted clinical decisions.
>
> **LLM regulatory void:** The FDA's existing framework was designed for narrowly scoped, deterministic AI/ML tools. Generative AI and foundation models, which can produce different outputs for identical inputs and are used across a wide range of clinical tasks, do not neatly fit into existing device classification categories.

Post-market surveillance limitations: While the FDA maintains adverse event reporting systems, the infrastructure for monitoring AI device performance drift, demographic performance disparities, and real-world effectiveness remains underdeveloped compared to the pace of deployment.

Ethical Implications

A number of ethical implications are important to consider with regard to AI in emergency medicine.

Bias Amplification at Scale

The ethical dimension of AI bias extends beyond individual algorithmic errors. When a biased algorithm is deployed across a health system serving millions of patients, the bias is amplified on a population scale. The Obermeyer algorithm affected care allocation for 200 million Americans [9]. If a biased LLM is integrated into a clinical decision support system used by thousands of physicians, every recommendation it generates could propagate systemic inequities.

This is not a hypothetical concern. The finding that 84 percent of clinical AI studies fail to report the racial composition of their training data [31] suggests that bias may be present in tools clinicians use daily, without any mechanism for detection. The combination of opaque algorithms, unreported demographics, and system-wide deployment creates conditions where bias can operate at unprecedented scale.

Equity: AI Working Worst for Those Who Need It Most

The ethical imperative is sharpened by a painful irony: AI tools tend to perform worse for the populations that face the greatest health disparities. Training datasets drawn predominantly from academic medical centers underrepresent rural, low-income, and minority patients. Algorithm performance validated on one demographic may degrade for another. And the digital divide means that the health systems serving the most vulnerable populations are the least likely to have the infrastructure, expertise, or resources to implement AI tools or detect when they are failing.

This creates a risk that AI becomes a force multiplier for existing inequities rather than a corrective. The promise of democratized expertise AI, bringing specialist-level diagnostics to underserved areas, is real but will remain unfulfilled unless deliberate policy, investment, and design choices prioritize equity.

ECRI's Warning: AI as the Number-One Health Technology Hazard

When ECRI, the independent, nonprofit patient safety organization that has published its annual health technology hazards list for decades, named AI the number-one health technology hazard for 2025, it was an institutional declaration that the risks of clinical AI had reached a threshold demanding immediate attention [10]. This was the first time AI had claimed the top spot, displacing more traditional hazards such as cybersecurity threats and device malfunctions.

ECRI's concerns centered on inadequate governance, insufficient validation, and the speed at which AI tools were being deployed relative to the maturity of safety frameworks. The designation served as a wake-up call for health system leaders: adopting clinical AI is not merely a technology decision. It is a patient safety decision.

Consent and Transparency

Patients have a right to know when AI is influencing their care, and current practice often falls short of that standard. Many patients are unaware that AI tools may be interpreting their imaging studies, flagging their vital signs for clinical review, or generating their physician's documentation. The absence of standardized consent and disclosure practices for clinical AI use raises fundamental questions about patient autonomy and the physician-patient relationship.

As AI becomes more deeply integrated into clinical workflows, the obligation to inform patients about its role becomes more urgent. This is not simply a legal question but an ethical one: the trust between patient and physician depends on transparency, and that trust is undermined when significant aspects of the care process are invisible.

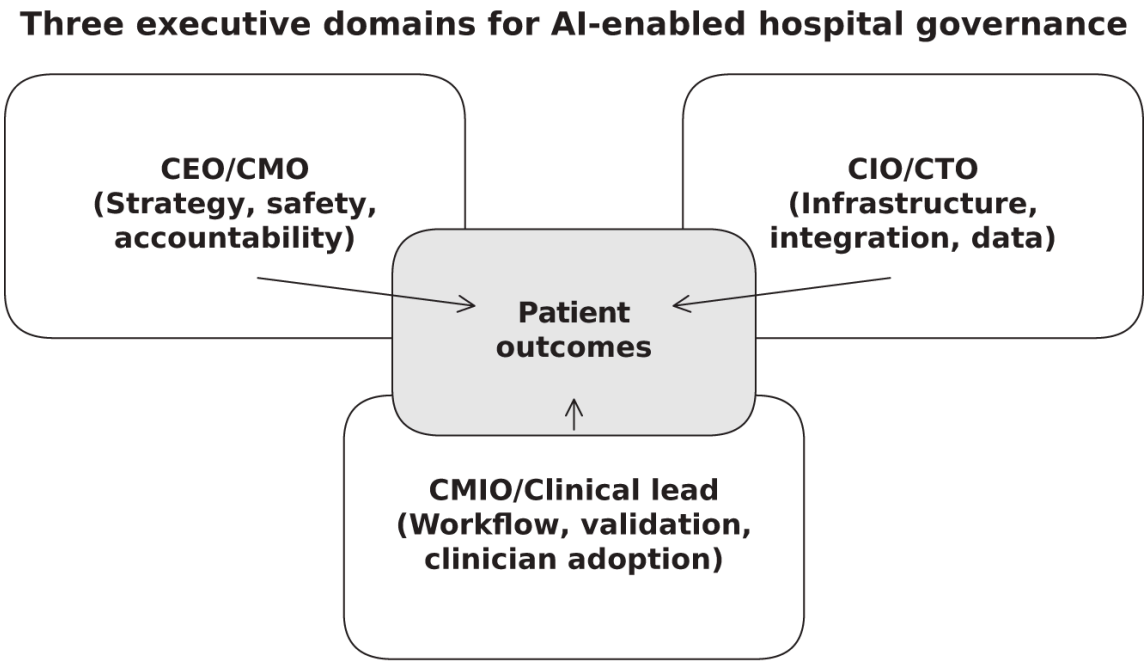
Leadership Implementation Framework

The successful integration of AI into healthcare requires coordinated leadership across three executive domains (see Figure 1.3). Each leader brings a distinct perspective, and the intersections between their responsibilities are where the most consequential decisions are made.

CEO: Strategic Vision and Governance

The chief executive officer (CEO) bears ultimate responsibility for the organization's AI strategy and must address three critical imperatives:

> **Strategic AI investment:** AI investment decisions must be grounded in organizational priorities, not vendor enthusiasm. The CEO should establish criteria for evaluating AI opportunities based on clinical impact, operational value, strategic alignment, and implementation feasibility. The total cost of ownership, including infrastructure, training, change management, and ongoing monitoring, often exceeds the technology acquisition cost by a factor of three to five.



Conceptual diagram based on the leadership implementation framework, Chapter 1.

Figure 1.3 The AI-enabled hospital ecosystem

Governance structure: Every health system deploying clinical AI needs a formal AI governance committee with representation from clinical leadership, information technology, legal, compliance, ethics, and patient advocacy. This committee should establish policies for AI evaluation, approval, deployment, monitoring, and decommissioning. The governance structure must be empowered to say no to AI tools that do not meet evidentiary standards, even when vendor pressure and competitive anxiety push toward adoption.

ROI measurement: AI investments must be measured against defined outcomes, not just financial returns, but clinical quality metrics, patient experience scores, clinician satisfaction, and equity indicators. The Kaiser TPMG ambient AI deployment provides a model: measurable outcomes (1,794 working days saved, 84 percent satisfaction) tied to organizational priorities (physician well-being, patient access) [27].

CIO: Infrastructure and Integration

The chief information officer (CIO) is the architect of AI readiness and must navigate three operational challenges:

Infrastructure readiness: AI deployment requires robust computational infrastructure, reliable network connectivity, and scalable data storage. Cloud-based AI tools depend on consistent, high-bandwidth connectivity that not all facilities can guarantee. The CIO must assess current infrastructure, identify gaps, and develop a phased investment plan that supports both current AI deployments and anticipated future growth.

EHR integration: Most clinical AI tools must integrate with the electronic health record to deliver value at the point of care. EHR integration is technically complex, operationally disruptive, and often more expensive than anticipated. The CIO should establish integration standards, invest in interoperability infrastructure (including FHIR-based APIs), and maintain a testing environment to validate new AI tools before clinical deployment.

Cybersecurity: AI tools that process protected health information introduce new attack surfaces and data governance challenges. The CIO must ensure that AI vendors meet organizational security standards, that data flows are encrypted and auditable, and that incident response plans account for AI-specific risks, including adversarial attacks, model poisoning, and data exfiltration.

CMO: Clinical Validation and Adoption

The chief medical officer (CMO) is the guardian of clinical quality and must champion three priorities:

Clinical validation: No AI tool should be deployed in clinical practice without rigorous validation in the specific population and workflow where it will be used. The CMO should establish validation protocols that include performance assessment across demographic subgroups, comparison to existing clinical processes, and prospective monitoring of clinical outcomes. The Epic Sepsis Model's performance gap between developer-reported metrics and independent validation [8] should serve as a standing reminder of why internal validation cannot substitute for independent assessment.

Physician champion programs: Successful AI adoption requires clinical champions—respected physicians who embrace the technology, model its use, and serve as peer educators. The CMO should identify and support these champions across specialties, providing them with protected time, training resources, and a direct channel to organizational leadership for feedback and concerns.

Workflow redesign: AI does not simply automate existing workflows; it changes them. The CMO must lead workflow redesign efforts that integrate AI tools into clinical processes to enhance, rather than disrupt, care delivery. This requires collaboration with frontline clinicians, iterative pilot testing, and willingness to modify or abandon approaches that do not improve the clinician or patient experience.

Future Outlook

The trajectory of AI in healthcare points toward several developments that, while not guaranteed, are grounded in current evidence and observable trends.

AlphaFold's Nobel Prize Moment and Drug Discovery

In October 2024, Demis Hassabis and John Jumper of Google DeepMind were awarded the Nobel Prize in Chemistry for their work on AlphaFold. This AI system solved the protein structure prediction problem, one of biology's grand challenges, for more than 50 years [37]. AlphaFold 3, published in *Nature* in May 2024, extended the system's capabilities to predict the structures of protein-nucleic acid complexes, small molecules, and post-translational modifications [37].

The numbers are staggering: AlphaFold has generated predicted structures for more than 240 million proteins, been accessed by more than 3.3 million researchers worldwide, and accumulated more than 40,000 citations [37]. Its impact on drug discovery is just beginning to be felt, but the potential is enormous. By accurately predicting protein folding and interactions, AlphaFold accelerates the identification of drug targets, the design of therapeutic molecules, and the understanding of disease mechanisms at the molecular level.

The Nobel recognition was significant not only for AlphaFold itself but for what it signaled about the scientific establishment's acceptance of AI as a transformative research tool. When the world's most prestigious scientific prize is awarded for AI-driven research, the message is unambiguous: AI is no longer an adjunct to science. It is a central instrument of discovery.

Ambient AI as Standard of Care

The evidence reviewed in this chapter, the TPMG deployment of 7,260 physicians [27], the multiple RCTs demonstrating burnout reduction [5,20,21], and the rapid adoption of Dragon Copilot across 600,000 clinicians [22] point toward a near future where ambient AI documentation is no longer optional but expected. The trajectory suggests that within three to five years, a health system that does not offer ambient documentation tools may struggle to recruit and retain physicians, particularly younger physicians who have trained with these tools.

The transition from innovation to standard of care carries implications for malpractice, accreditation, and quality measurement. If ambient AI documentation becomes the norm, will failure to implement it constitute a deviation from the standard of care? The legal and regulatory frameworks have not yet addressed this question, but the clinical momentum is unmistakable.

Regulatory Maturation

The next several years will likely see the regulatory landscape mature from its current patchwork state toward more comprehensive frameworks. The EU AI Act's phased implementation through 2027 will create the world's first comprehensive enforcement regime for medical AI [34]. In the United States, the FDA's PCCP framework [13] will be tested and refined as more devices incorporate predetermined change control plans. The growing influence of CHAI and the Joint Commission [36] will push health systems toward formal AI governance structures.

The most consequential regulatory question remains unanswered: who is liable when AI contributes to a medical error? Until legislatures and courts establish clear frameworks, the uncertainty will continue to shape how

health systems, physicians, and AI developers approach deployment and oversight.

Multimodal AI and the Integration Frontier

Looking toward 2026 and beyond, the most significant technical development may be the emergence of multimodal AI systems that integrate imaging, genomics, electronic health record data, and real-time physiological monitoring into unified predictive frameworks. Early-stage research suggests that these systems may achieve diagnostic and prognostic accuracy that exceeds that of any single data modality.

The integration of AI across data types and care settings from primary care to emergency medicine to post-acute care represents the next frontier. The successes of 2024 and 2025 were largely domain-specific: AI excelling in radiology, in documentation, in sepsis prediction. The future points toward AI systems that function across these domains, creating a more comprehensive and continuous picture of patient health.

Key Takeaways

AI is operational, not aspirational. With 1,250+ FDA-cleared devices, 66 percent physician adoption, and a $15 billion global market, healthcare AI has moved from research to routine practice.

The evidence base for landmark applications is strong. The MASAI trial (29 percent more cancers detected), ambient AI RCTs (30 percent burnout reduction), and sepsis AI (17 percent mortality reduction) demonstrate genuine clinical value supported by rigorous study designs.

High-profile failures contain essential lessons. IBM Watson’s $4 billion collapse and Epic’s sepsis model controversy reveal the dangers of inadequate validation, proprietary opacity, and the gap between marketing claims and real-world performance.

Algorithmic bias is a patient safety crisis. Racial, gender, and socioeconomic biases are pervasive in healthcare AI, with 84 percent

of studies failing to report training data demographics and 90.9 percent of LLM studies demonstrating racial bias.

Regulatory frameworks are maturing but incomplete. The FDA PCCP guidance, EU AI Act, and Joint Commission/CHAI guidance represent progress, but no federal algorithmic audit mandate or AI liability framework exists.

The digital divide threatens to widen health disparities. AI tools work worst for the populations that need them most, and rural and underserved hospitals are least likely to adopt them.

Leadership must be deliberate, not reactive. Successful AI implementation requires coordinated leadership from the CEO, CIO, and CMO, with formal governance, rigorous validation, and ongoing monitoring, not just technology procurement.

References

1. Bipartisan Policy Center. FDA oversight: understanding the regulation of health AI tools. Washington (DC): Bipartisan Policy Center; 2025 Nov. Available from: https://bipartisanpolicy.org/issue-brief/fda-oversight-understanding-the-regulation-of-health-ai-tools.

2. American Medical Association. AMA physician survey on AI in healthcare [Internet]. Chicago (IL): American Medical Association; 2025 Feb.

3. MarketsandMarkets. Artificial intelligence in healthcare market global forecast to 2030 [Internet]. 2024.

4. Lång K, Josefsson V, Larsson AM, Larsson S, Högberg C, Sartor H, et al. Artificial intelligence-supported screen reading versus standard double reading in the Mammography Screening with Artificial Intelligence trial (MASAI): a clinical safety analysis of a randomised, controlled, non-inferiority, single-blinded, screening accuracy study. Lancet Oncol. 2023 Aug; 24(8):936-944. doi:10.1016/S1470-2045(23)00298-X. PMID: 37541274.

5. Lukac PJ, Turner W, Vangala S, Chin AT, Khalili J, Shih YT, et al. Ambient AI scribes in clinical practice: a randomized trial. NEJM AI. 2025. PMID: 40672471. DOI: 10.1101/2025.07.10.25331333 (medRxiv preprint 11 Jul 2025; NEJM AI announcement 18 Dec 2025).

6. Adams R, Henry KE, Sridharan A, Soleimani H, Zhan A, Rawat N, et al. Prospective, multi-site study of patient outcomes after implementation of the TREWS machine learning-based early warning system for sepsis. Nat Med. 2022; 28(7):1455-1460. PMID: 35864252. DOI: 10.1038/s41591-022-01894-0.

7. Ross C, Swetlitz I. IBM's Watson supercomputer recommended 'unsafe and incorrect' cancer treatments, internal documents show [Internet]. STAT News. 2018 July.

8. Wong A, Otles E, Donnelly JP, Krumm A, McCullough J, DeTroyer-Cooley O, et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Intern Med. 2021. doi:10.1001/jamainternmed.2021.2626.

9. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019 Oct; 366(6464):447-453. doi:10.1126/science.aax2342.

10. ECRI. Top 10 Health Technology Hazards for 2025 [Internet]. Plymouth Meeting (PA): ECRI; 2024 Dec.

11. Patel TA, et al. Machine learning-enabled medical devices authorized by the US FDA in 2024: a cross-sectional analysis. Biomedicines. 2025 Dec. PMID: PMC12730494. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12730494.

12. Innolitics. 2025 year in review: AI/ML medical device 510(k) clearances [Internet]. 2025 Dec. Available from: https://innolitics.com/articles/year-in-review-ai-ml-medical-device-k-clearances.

13. U.S. Food and Drug Administration. Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions. Final guidance

[Internet]. Silver Spring (MD): FDA; 2024 Dec 4. Available from: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial-intelligence.

14. Rock Health. 2025 digital health funding report [Internet]. 2025.

15. Office of the National Coordinator for Health Information Technology (ONC/ASTP). Hospital use of predictive AI [Internet]. 2024.

16. Naik N, et al. Clinician perspectives on de-skilling risks of AI in healthcare. JMIRx Med. 2024.

17. Yao J, Zhou Z, Cui H, Ouyang Y, Han W. Trust transfer from medical AI to doctors and hospitals: integrating digital, AI, and scientific literacy in a cross-sectional framework. BMC Med Ethics. 2025 Oct 17. PMID: 41107871. DOI: 10.1186/s12910-025-01300-7.

18. Brainomix. AI-powered stroke imaging: National Health Service outcomes. Lancet Digital Health. 2025 Dec.

19. Shlobin NA, et al. Impact of Viz.ai on stroke care: a systematic review and meta-analysis. Neurology. 2025 Apr. doi:10.1212/WNL.0000000000208383.

20. Providence Health. Ambient AI documentation and physician burnout: a randomized controlled trial. Future Healthc J. 2025 July.

21. Multicenter quality improvement study. Ambient AI and physician burnout reduction. JAMA. 2025.

22. Nuance Communications/Microsoft. Dragon Copilot adoption metrics [Internet]. 2025.

23. Boussina AE, et al. COMPOSER: an AI-based sepsis prediction system. npj Digital Med. 2024 Jan.

24. HERACLES Study Group. AI-driven sepsis detection across 97,559 hospital stays. npj Digital Med. 2026 Jan.

25. JAMA Health Policy Summit. AI implementation in healthcare: summit findings. JAMA. 2025 Oct.

26. Lee B, Kramer P, Sandri S, Chanda R, Favorito C, Nasef O, et al. Early recalls and clinical validation gaps in artificial intelligence-enabled medical devices. JAMA Health Forum. 2025; 6(8):e253172. PMID: 40844774. DOI: 10.1001/jamahealthforum.2025.3172.

27. The Permanente Medical Group, Kaiser Permanente. Ambient AI documentation deployment at scale. NEJM Catal. 2025.

28. Sjoding MW, Dickson RP, Iwashyna TJ, Gay SE, Valley TS. Racial bias in pulse oximetry measurement. N Engl J Med. 2020; 383:2477-2478.

29. Systematic review of LLM bias across 24 studies. 2025. Cited in: Yale Global Health Review. 2026 Jan.

30. Bouguettaya A, Stuart EM, Aboujaoude E. Racial bias in AI-mediated psychiatric diagnosis and treatment. npj Digit Med. 2025 Jun 4. PMID: 40467886. DOI: 10.1038/s41746-025-01746-4.

31. Systematic review of 390 clinical AI studies. Demographic transparency in clinical AI/ML research. Oxford Open Digital Health. 2025.

32. Nature Health. AI adoption disparities in rural and underserved hospitals. 2026 Jan.

33. Federal Communications Commission. Broadband deployment report [Internet]. Available from: https://www.fcc.gov/reports-research/reports/broadband-progress-reports.

34. European Union. Regulation (EU) 2024/1689 (Artificial Intelligence Act). Off J Eur Union. 2024 Aug 1.

35. World Health Organization. Ethics and governance of artificial intelligence for health: guidance on large multimodal models Internet]. 2024 Jan. Available from: https://www.who.int/publications.

36. Joint Commission/Coalition for Health AI (CHAI). AI governance guidance for accredited healthcare organizations [Internet]. 2025 Sep.

37. Abramson J, Adler J, Dunger J, Evans R, Green T, Pritzel A, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. Nature. 2024 May.

AI in Emergency Medicine: What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution, First Edition. Harvey Castro.
© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

2
Real Patients, Real AI: From Burnout to Breakthroughs

I had worked my entire life to become a doctor. And the night my body tried to kill me, I had just gotten there.

I had graduated emergency medicine residency. I was finally an attending, my first real shift in the ER, the title I had bled for. Then the migraines started. These were not the garden-variety tension headaches you shake off with Ibuprofen. These were blinding, nauseating, something-is-wrong migraines that sent me, a man who avoids doctors the way his mother taught him, to seek help.

The doctors ran everything. X-rays. CT scans. MRIs. A lumbar puncture, a needle between the vertebrae of my spine, one of the most painful procedures in medicine. I know, because I've ordered it on patients. Having it done to me was a different education entirely. After all of it, the diagnosis was a shrug. "Harvey, here's a pill. The next time you get a headache, take this. You'll be fine."

I took the pill.

Thirty minutes later, I turned to my wife and said three words: "Call 911. Now."

By the time the paramedics arrived, they checked my blood sugar. It had dropped to 20. I had just eaten dinner. A blood sugar of 20 means your brain is starving. It means you are minutes from death. I looked at the paramedic and said, "Please put a tube down my throat. I'm about to die in front of you." Then I blacked out.

When I came to, I was in the emergency room. My colleagues, the same doctors I had just started working alongside, surrounded the gurney. They were asking me basic identification questions. My date of birth. My name. Standard protocol. I know this protocol. I've run this protocol.

I couldn’t answer. My entire left side was paralyzed. When I tried to speak, the only sound that came out was: one, one, one. Over and over. One, one, one. I was a doctor trapped in his own body, able to hear everything—the beeping monitors, the worried whispers, the overhead pages—but unable to say a single coherent sentence.

I spent seven days in the ICU. Seven days of consciousness drifting in and out, of hearing nurses discuss my case in the hallway, of understanding every word and being unable to respond. When I finally emerged, I spoke like a stroke patient, halting, searching for words that used to come as naturally as breathing.

It took me months of my own research to piece together what had happened. The pill they gave me for migraines had triggered an adverse drug reaction. The reaction caused a vasospasm in my brain, the blood vessels clamped down, mimicking a stroke. No one caught it because no one was looking for it. The imaging had been done. The labs had been run. The box had been checked. But the interaction between that drug and my physiology slipped through every crack in the system.

I wish I had AI. I wish I had something, a monitoring system, a clinical decision tool, anything, to help me understand what was happening to my body before it nearly killed me.

Executive Summary

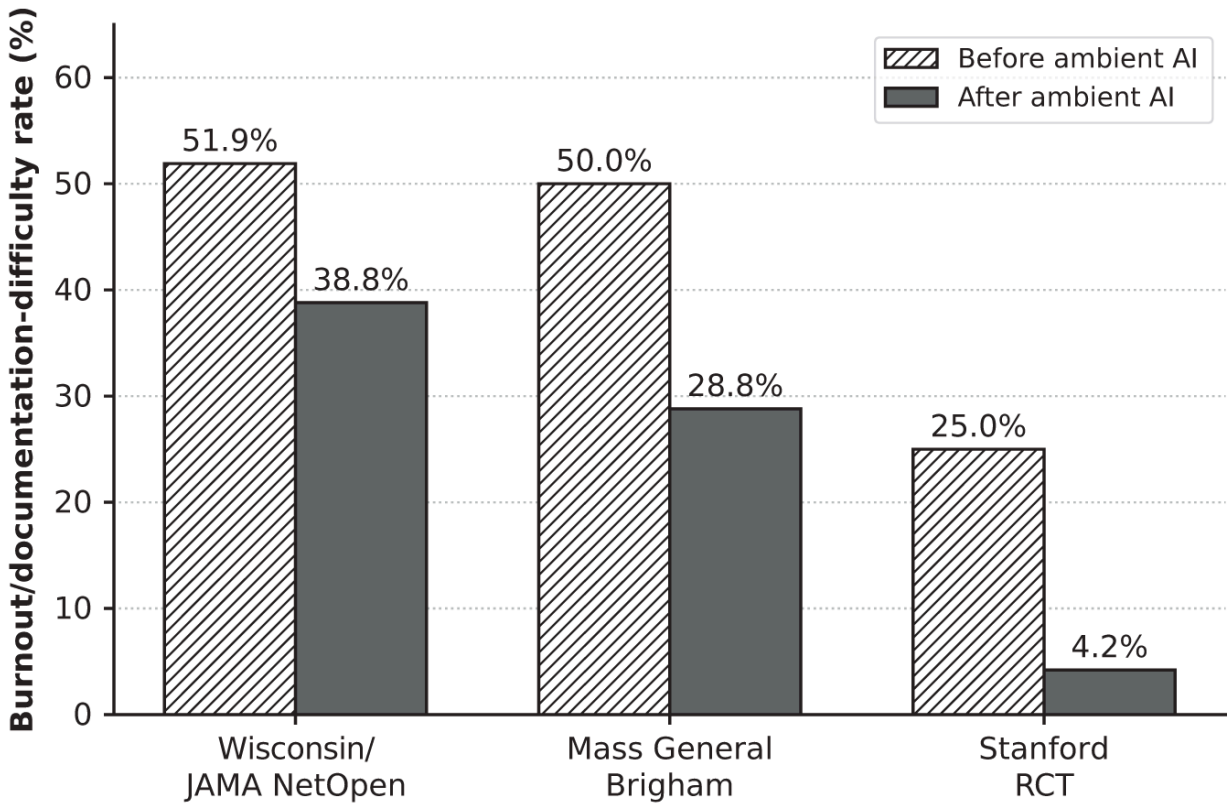
A drawing has haunted me since the day I first saw it. A 7-year-old girl sketched her visit to the doctor: herself on the examining table, small and alone, and the physician, with his back turned to her, staring at a computer screen. The physician published that image in the *Journal of the American Medical Association* in 2012, alongside a question that still echoes: “The economic stimulus bill has directed $20 billion to healthcare information technology, largely funding electronic medical record incentives. I wonder how much this technology will really cost?”

As an emergency physician, I know the cost. It is measured not only in dollars but in the erosion of the very thing that drew most of us to medicine, the human connection between a patient in distress and the clinician who reaches out a hand. The electronic health record is strangling that

connection. For every hour a physician spends face-to-face with a patient, nearly two additional hours are spent on documentation and desk work [1]. Physicians log into the EHR at 10 p.m. to finish notes they could not complete during the day, sacrificing evenings with their families in what Christine Sinsky, MD, who first measured this burden, memorably christened “pajama time” [2]. Nearly half of all American physicians report symptoms of burnout, and the consequences ripple outward: burned-out clinicians are more than twice as likely to make a major medical error, and an estimated 300 to 400 physicians die by suicide every year [3–5].

This chapter is about what happens when that arc begins to bend. A new generation of ambient artificial intelligence tools that listen to the physician-patient conversation, generate a clinical note, and let the doctor close the laptop is demonstrating, in randomized controlled trials and real-world deployments spanning millions of encounters, that the documentation burden can be meaningfully reduced. At Kaiser Permanente, 7,260 physicians used ambient AI across more than 2.5 million patient encounters, saving an estimated 15,791 hours of documentation time, and nearly half of the patients surveyed said their doctor spent less time looking at a computer [6]. At the University of Wisconsin, the first stepped-wedge randomized trial of ambient AI found that the proportion of practitioners exceeding the burnout threshold fell from 56.1 percent to 35.4 percent (see Figure 2.1) [7].

But this chapter is also about what lies beyond documentation. AI is detecting sepsis hours before the best clinicians, flagging cancers missed by human eyes, and shortening the diagnostic odyssey for patients with rare diseases from a decade to minutes. The breakthroughs are real. So are the risks. What follows is the evidence and the human stories it cannot fully capture.



Sources: Olson et al., JAMA Network Open (Oct 2025); Mishuris et al., JAMA Network Open (Aug 2025); Stanford RCT (manuscript ref. 21).

Figure 2.1 Physician time allocation before and after ambient AI documentation

Current Evidence

The evidence base for physician burnout and the tools designed to address it has grown substantially in recent years. What follows is a synthesis of the most rigorous data available, drawn from national surveys, randomized trials, and large-scale deployments.

The Burnout Crisis by the Numbers

The data on physician burnout tells a story that is both alarming and, in its most recent chapters, cautiously hopeful. The longest-running longitudinal series conducted by Tait Shanafelt, MD, and colleagues at Stanford, Mayo Clinic, the AMA, and the University of Colorado has tracked burnout at consistent intervals since 2011. In their most recent publication in *Mayo*

Clinic Proceedings (April 2025), based on surveys of 7,643 physicians conducted in late 2023 and early 2024, the burnout rate stood at 45.2 percent a steep decline from the all-time high of 62.8 percent recorded during the Omicron wave in 2021, but essentially unchanged from the 45.5 percent baseline when the series began more than a decade ago [8]. (Most burnout data derive from validated self-report instruments, chiefly the Maslach Burnout Inventory and the Mini-Z survey, administered periodically to large physician cohorts.)

The trajectory is revealing. Burnout rose from 45.5 percent in 2011 to 54.4 percent in 2014, coinciding with the surge in EHR adoption driven by the HITECH Act and Meaningful Use requirements. It dipped to 43.9 percent in 2017 as systems adapted and then fell further to a surprising low of 38.2 percent in 2020—a paradoxical “rally effect” during the early pandemic months when patient volumes dropped and the nation applauded healthcare workers. Then came the Omicron wave, and burnout exploded to 62.8 percent in 2021, the highest rate ever recorded. The subsequent decline to 45.2 percent is real, but the landing spot is essentially where we started.

That plateau is, by itself, a damning indictment. Despite more than a decade of wellness programs, mindfulness retreats, and institutional task forces, nearly half of all American physicians still meet the threshold for burnout. And the relative burden is worsening: physicians are now 82.3 percent more likely to experience burnout than other U.S. workers, after adjusting for age, gender, relationship status, and work hours [8]. Female physicians face a 27 percent greater risk than their male counterparts [8]. Emergency medicine, my own specialty, has the highest burnout rate of any field, at 63 percent [9].

Separate data streams corroborate the picture. The Medscape Physician Burnout & Depression Report for 2024, surveying 9,226 physicians, found an overall burnout rate of 49 percent, with 62 percent naming bureaucratic tasks as the primary driver [10]. The AMA’s Organizational Biopsy survey of roughly 18,000 physicians across 43 states reported burnout at 43.2 percent in 2024, the lowest in their series since COVID, but cautioned that job stress remained elevated at 45.1 percent [11]. The Physicians Foundation’s 2025 Well-being Survey found that 54 percent of physicians reported feelings of burnout. Despite a modest decline from 60 percent the prior year, the foundation noted that “dramatic shifts in 2025 have driven

stress and anxiety among physicians to levels not seen since the height of the pandemic" [12].

Nursing is no less affected. The American Nurses Foundation's 2023 Pulse Survey of more than 7,400 nurses found that 56 percent experienced burnout, 64 percent reported "a great deal of stress because of their job," and 39 percent intended to leave their current position within six months [47]. The NCSBN's 2024 National Nursing Workforce Study, surveying 800,000 nurses, found that approximately 138,000 nurses (RNs and LPN/VNs combined) had left the workforce since 2022, and nearly 40 percent intended to leave or retire within five years [48]. The average cost of replacing a single staff registered nurse (RN) now stands at $61,110, and the average acute-care hospital lost $4.75 million to RN turnover in 2024 alone [49].

The workforce consequences among physicians are equally severe. A JAMA Network Open study of 18,719 academic physicians found that 32.6 percent had moderate or greater intent to leave their institution within two years, and burned-out physicians were twice as likely to actually depart [13]. The AAMC projects a physician shortage of between 13,500 and 86,000 (the wide range reflects methodological differences: some estimates count only physicians who leave medicine entirely, while others include those who cut hours, shift to nonclinical roles, or retire early) by 2036, and if underserved populations could access care at the same rate as better-served populations, the shortfall would be approximately 202,800 [14]. The financial toll is staggering: an economic analysis by Han, Shanafelt, Sinsky, and colleagues published in *Annals of Internal Medicine* estimated that physician burnout costs the U.S. healthcare system $4.6 billion annually through turnover and reduced clinical hours, approximately $7,600 per physician per year [15]. Replacing a single physician costs an estimated $500,000 to $1,000,000, including lost billings, recruitment, and onboarding [50]. The Commonwealth Fund reported in 2024 that 39 percent of burned-out primary care physicians planned to stop seeing patients within one to three years, compared with 21 percent of those without burnout, and that 86 percent of burned-out physicians were dissatisfied with the time available per patient [51]. As Shanafelt observed in the most recent Stanford Medicine news release: "Many physicians still love what they do, but they just can't keep doing it at this pace in the current practice

environment, with its administrative burdens and regulatory burdens, and the proliferation of asynchronous messaging with patients through the electronic health record" [8].

The Documentation Burden: Where It All Begins

The single most consistent driver of burnout across surveys and specialties is documentation. Christine Sinsky's landmark 2016 time-and-motion study in *Annals of Internal Medicine*, the foundational research that reshaped how we think about physician workload, directly observed 57 physicians across four specialties and found that they spent 49.2 percent of their office day on EHR and desk work, compared with just 27.0 percent on direct clinical face time [1]. For every hour spent with a patient, nearly two additional hours were spent in the EHR, plus an average of an hour-and-a-half of personal time each evening [1].

The phrase Sinsky coined for this burden, "Other physicians and I began to wonder if we were spending more time caring for the computer than caring for the patient," became a rallying cry for a generation of clinicians who felt that the promise of health IT had been betrayed by its implementation [2].

Sinsky's work was not a one-time snapshot. Her 2024 follow-up study in the *Journal of General Internal Medicine*, drawing on Epic Signal data from 200,081 physicians at 396 organizations, found that primary-care physicians were spending 7.3 hours in the EHR for every 8 hours of scheduled patient time, with 2.7 of those hours occurring outside of clinic hours, in what Sinsky called "uncompensated labor" [16]. A separate longitudinal study in *Annals of Family Medicine*, tracking 141 academic primary care physicians from 2019 to 2023, documented a 7.8 percent increase in total EHR time per scheduled session, with order time surging 58.9 percent and patient message volume rising 55.5 percent above pre-pandemic levels [17]. The trend was clear: even as the pandemic receded, the EHR burden intensified.

And so the nightly ritual continues. Sinsky's research on family physicians found that they spent an average of 86 minutes per night on the EHR after office hours—the original "pajama time" measurement [2]. A 2024 study in *Annals of Family Medicine* found that one-third of upper-year family medicine residents spent three or more hours per night on after-hours EHR

work. That high pajama time was associated with lower medical knowledge scores, lower professional satisfaction, and higher burnout [18].

The human cost of that documentation burden is not abstract. Burned-out physicians are 2.2 times more likely to report a major medical error, according to Tawfik, Shanafelt, and colleagues in a landmark 2018 study of 6,695 U.S. physicians in *Mayo Clinic Proceedings*, and physicians working in the lowest-graded safety environments (hospitals graded C, D, or F by the Leapfrog Group, based on rates of medical errors, infections, and other preventable harms) were approximately three times more likely to report errors than those in the highest-graded environments (OR ~3.12) [3]. A 2024 meta-analysis in *JAMA Network Open*, analyzing 85 studies and 288,581 nurses across 32 countries, confirmed that nurse burnout is associated with lower patient safety, lower patient satisfaction, and lower quality of care, and that the association has grown more negative over three decades [19]. As Shanafelt himself observed, “Both burnout and medical errors independently double the risk of suicidal thoughts among physicians” [4].

The Emergence of Ambient AI: Evidence from Randomized Trials

Against this backdrop of crisis, a new body of rigorous evidence has emerged that ambient AI documentation tools can meaningfully reduce the burden. In November 2025, three landmark randomized controlled trials were published simultaneously in *NEJM AI*, marking the first time ambient AI had been subjected to the gold standard of clinical evidence.

The UCLA Trial Lukac, Mafi, and colleagues conducted a parallel three-group pragmatic randomized controlled trial (RCT) at UCLA Health, randomizing 238 outpatient physicians across 14 specialties to Microsoft DAX Copilot, Nabla, or usual care. Over 72,349 outpatient visits, Nabla produced a statistically significant 9.5 percent reduction in time-in-note compared to control. DAX showed a nonsignificant 1.7 percent reduction. Critically, both tools produced significant improvements in burnout-related secondary outcomes: the Mini-Z composite score improved by 2.83 points in the DAX arm and 2.69 in the Nabla arm, and work exhaustion decreased

in both arms [20]. (The Mini-Z is a validated single-item burnout screening tool used to quickly assess physician burnout and satisfaction.)

The UW Health Trial Afshar and colleagues conducted a 24-week stepped-wedge individually randomized trial across ambulatory clinics in Wisconsin and Illinois, enrolling 66 practitioners and analyzing 71,487 clinical notes. Using Abridge, they found a 0.44-point reduction in work exhaustion on the Stanford Professional Fulfillment Index ($P < 0.001$), the primary outcome. The proportion of practitioners exceeding the burnout threshold fell from 56.1 percent to 35.4 percent. Daily time on notes decreased by 21.6 minutes ($P < 0.001$). The number needed to treat was 1.68, meaning that for every 1.68 practitioners given access to the tool, one achieved a clinically meaningful reduction in burnout. Remarkably, AI-generated notes also showed higher ICD-10 billing compliance than human-authored notes (6.87 versus 5.94 on a 10-point scale; $P < 0.001$) [7].

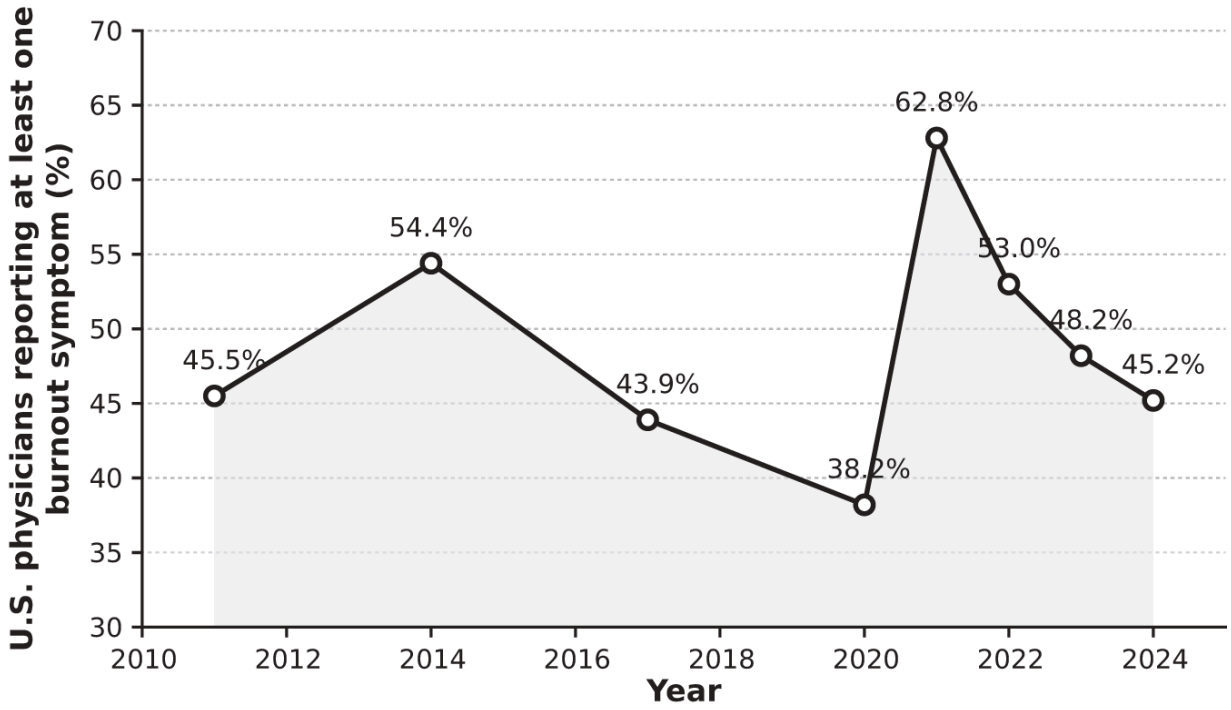
The Providence Trial In a smaller but important RCT published in *Future Healthcare Journal*, Wendt and colleagues at Providence Health randomized 24 family medicine physicians, selected based on Epic Signal data showing the highest documentation burden, to DAX Copilot. The results were striking: a 30.3 percent reduction in burnout, a 2.5-hour weekly reduction in documentation burden, a 49.5 percent decrease in documentation frustration, and a near-elimination of difficulty engaging with patients (from 25 percent of providers reporting difficulty to just 4.2 percent) [21].

Large-Scale Validation: JAMA and Kaiser

Beyond RCTs, two large-scale studies provided crucial real-world validation. The JAMA Network Open multicenter quality improvement study, published in October 2025 by Olson, Schwamm, and colleagues across six academic and community health systems, enrolled 263 ambulatory clinicians and measured outcomes at 30 days. Self-reported burnout dropped from 51.9 percent to 38.8 percent, which is an adjusted odds ratio of 0.26. Note-related cognitive task load fell by 2.64 points on a 10-point scale ($P < 0.001$), and after-hours documentation decreased significantly [22].

The Kaiser Permanente deployment, published by Tierney, Liu, Lee, and colleagues in *NEJM Catalyst* in March 2025, represents the largest real-world evaluation of ambient AI in healthcare to date. Over 63 weeks, 7,260 TPMG physicians used ambient AI across 2,576,627 patient encounters, saving an estimated 15,791 hours of documentation time. Among surveyed physicians, 88 percent reported a positive impact on visit interactions, and 67 percent used the technology five or more days per week [6]. The sheer scale of thousands of physicians and millions of encounters transformed ambient AI from a promising pilot into an operational reality.

Additional evidence continued to accumulate across the country. At Mass General Brigham, a study of 1,430 clinicians across two academic medical centers published in *JAMA Network Open* in August 2025 found that ambient documentation was associated with an absolute reduction of 21.2 percentage point (the raw percentage point drop, as opposed to relative change) in burnout prevalence at 84 days. Rebecca Mishuris, MD, the chief medical information officer at Mass General Brigham, stated: “There is literally no other intervention in our field that impacts burnout to this extent” [54]. At the University of Chicago, a matched-cohort analysis found that ambient AI users spent 8.5 percent less total time in the EHR and more than 15 percent less time composing notes than matched controls [55]. At the University of Pennsylvania, a JAMA Network Open study found a 20.4 percent reduction in time per note, a 30 percent reduction in after-hours documentation, and a 9.3-percentage-point improvement in same-day chart closure [56].



Source: Shanafelt et al., Mayo Clinic Proceedings (multiple years); AMA Physician Burnout & Well-being reports.

Figure 2.2 Physician burnout rate trends, 2011–2024

Taken together, the evidence base for ambient AI now includes three RCTs, at least four multicenter or large-scale studies, and multiple single-center quality-improvement analyses, an unusually robust body of evidence for any healthcare technology, let alone one that reached widespread deployment less than three years after its initial introduction (see Figure 2.2).

Real-World Cases: The Kaiser Permanente Transformation

When Brian Hoberman, MD, the executive vice president of IT and chief information officer at the Permanente Federation describes what happened when Kaiser Permanente deployed ambient AI across its physician workforce, he reaches for a phrase that would be hyperbolic in almost any other context. “It feels like magic,” he told the American Medical Association, “to the docs, and often to the patients” [23].

For a system that has invested for decades in health information technology, Kaiser Permanente was an early pioneer of comprehensive EHR adoption. The ambient AI rollout represented something qualitatively different from prior technology deployments. The EHR had delivered measurable gains in safety, quality, and coordination. But it had also exacted a toll that clinicians felt every day in their exam rooms and every evening in their homes. “The ultimate goal of healthcare IT was always: make me a better doctor,” Hoberman explained in a Becker’s Hospital Review interview at HIMSS 2025. “We did a lot of things that made the healthcare delivery system much more effective in delivering quality and a lot of other great things. But people didn’t feel it, and then you come along with ambient, and the clinicians really feel it, and their patients feel it” [24].

The TPMG deployment began in October 2023, initially as a carefully controlled pilot with more than 1,000 clinicians across 63,000 patient encounters. The pilot included a rigorous quality assurance process: researchers scored AI-generated transcripts using standardized rubrics, averaging 48 out of 50 across measures of accuracy, bias, consistency, and succinctness [25]. Only after demonstrating that quality threshold did the system proceed to enterprise rollout across all eight Kaiser Permanente regions, encompassing 600 medical offices and 40 hospitals.

By December 2024, 63 weeks after initial deployment, the numbers were remarkable. A total of 7,260 physicians had used the technology, with 3,447 having used it in 100 or more patient encounters each. The total time saved in documentation was estimated at 15,791 hours, roughly 1,974 eight-hour workdays. The specialties with the highest adoption rates were those with the heaviest documentation loads: mental health, emergency medicine, primary care, allergy, and cardiology [6].

But the statistics that matter most may be those that capture the human experience. In a survey of 118 patients, 47 percent reported that their physician spent less time at the computer during the visit. More than one-third said their doctor spent more time speaking directly to them. Fifty-six percent reported a positive impact on the quality of their visit, and no patient reported a negative impact. Two-thirds were comfortable with the technology [6].

Hoberman himself described the change in terms of what ambient AI restored—not what it added. “As you and I are talking right now, I’m looking at you, and you’re looking at me,” he told Becker’s. “That’s what ambient does in the exam room. So it makes the doc and the patient, or the practitioner and the patient, connect more” [24]. The technology, he noted, was also “a better listener than human memory.” Physicians reviewing their AI-generated notes would sometimes discover clinically important details that the conversation had surfaced but that they had not consciously registered [24].

Vincent Liu, MD, the chief data officer at TPMG and a co-author on the *NEJM Catalyst* study, framed the finding in institutional terms: “We have now shown that this technology alleviates workloads for doctors. Both doctors and patients highly value face-to-face contact during a visit, and the AI scribe supports that” [26]. Kaiser Permanente subsequently expanded the technology to more than 25,000 clinicians, system-wide physicians, advanced practice providers, and pharmacists, making it one of the largest AI deployments in American healthcare.

Ramin Davidoff, MD, the executive medical director of the Southern California Permanente Medical Group, articulated the institutional philosophy that guided the deployment: “Creating space for the patient and the physician connection is what inspired us to implement this technology” [52]. That framing technology in service of human connection, rather than in opposition to it, is what distinguished the Kaiser deployment from so many prior health IT rollouts that promised efficiency but delivered frustration.

The Kaiser story illustrates what it looks like when an institution commits to technology not as a productivity lever but as a means of restoring something lost. The physicians did not describe ambient AI as a tool. They described it as magic, the magic of being allowed, at last, to be a doctor in a room with a patient simply.

Risks and Limitations

The enthusiasm for ambient AI is genuine and evidence-based, but it must be tempered by an honest accounting of what these tools cannot yet do,

what they sometimes get wrong, and what dangers they may introduce.

Note Quality and Hallucinations

The most technically concerning finding in the current evidence is that AI-generated clinical notes contain *hallucinations*, fabricated or inaccurate content at measurable rates. A study published in *Frontiers in Artificial Intelligence* comparing Suki ambient AI notes to physician-authored "gold standard" notes found hallucinations in 31 percent of AI-generated notes, compared with 20 percent of human-authored notes [27]. While physician-authored notes scored modestly higher overall (4.25 versus 4.20 on the PDQI-9 quality scale), ambient notes were rated more thorough and better organized, suggesting the errors were of a different character than omissions in rushed human notes. The UCLA RCT found that physicians rated inaccuracies as occurring "occasionally," with the most common types being omissions, structural concerns, and pronoun resolution errors (instances where the AI misidentifies who a pronoun such as "he" or "she" refers to, potentially attributing symptoms to the wrong person) [20]. At Kaiser, hallucinations were described as "quite rare" in quality audits. Still, the fact that a clinical note can contain fabricated content that a busy physician might miss raises a safety question the field has not yet fully resolved [28]. The apparent contradiction dissolves on closer inspection: Kaiser's internal audits found hallucinations rare in routine encounters, but the broader research, which includes more complex cases and different clinical settings, shows they occur at measurable rates. Both findings reinforce the need for ongoing monitoring.

Robert Wachter, MD, the influential chair of medicine at UCSF, articulated the deepest concern: "Humans stink at maintaining vigilance over time" [28]. As physicians grow accustomed to ambient AI and stop reading every line of every note with the same scrutiny they applied in the first week, the risk of undetected errors may increase, not decrease. This "de-skilling" or "automation complacency" risk is well documented in aviation and other high-reliability fields, and healthcare has no reason to believe it is immune.

The Systemic Nature of Burnout

It would be a dangerous oversimplification to conclude that ambient AI solves the burnout crisis. Burnout is, as the AMA and the National

Academy of Medicine have repeatedly emphasized, a system issue, not a resiliency deficit [11]. The drivers of physician distress extend far beyond documentation to include inadequate staffing, misaligned compensation structures, loss of autonomy, moral injury from insurance barriers, and a culture that stigmatizes mental health help. The Physicians Foundation found that 73 percent of physicians agree that stigma still exists around seeking mental health support [12]. Ambient AI addresses one admittedly major contributor to burnout. Still, it cannot fix broken systems, restore gutted nursing ratios, or repair the fractured relationship between healthcare organizations and their clinicians.

The evidence itself reflects this limitation. In the Sutter Health study published in *JAMA Network Open*, burnout declined from 42.1 percent to 35.1 percent after implementation. Still, the change was not statistically significant, a reminder that even substantial-seeming improvements may fall short of statistical thresholds in the real world [29]. The JAMA multicenter study showed burnout falling from 51.9 percent to 38.8 percent, which is clinically meaningful but still leaves more than a third of clinicians burned out [22].

The Gap Between RCTs and Real-World Deployment

The three ambient AI RCTs were conducted under carefully controlled conditions at well-resourced academic health systems with dedicated implementation teams, physician informaticists, and structured training programs. The utilization rates tell part of the story: at UW Health, utilization reached 71 percent, but at UCLA, DAX was used in only 33.5 percent of eligible visits and Nabla in 29.5 percent [20]. In the messy reality of community hospitals, safety-net systems, and under-resourced clinics, adoption rates may be substantially lower, and the gap between demonstrated efficacy and realized effectiveness may be wide. The Stanford qualitative follow-up found that physicians wanted notes tailored to their personal documentation style, that procedural and templated specialties gained less benefit, and that early-version AI notes were sometimes wordy or used unnatural language friction points that do not appear in aggregate outcome statistics but that determine whether a physician uses the tool on Monday morning or abandons it by Wednesday [53].

Alert Fatigue and Decision-Support Overload

Beyond documentation, AI tools in clinical settings face the persistent challenge of alert fatigue. Clinicians in EHR-integrated environments are already bombarded with warnings, flags, and pop-ups. Adding AI-driven clinical decision support for sepsis, for medication interactions, and for diagnostic suggestions risks compounding the very cognitive overload that ambient documentation tools aim to relieve. The balance between timely clinical alerts and information paralysis remains unresolved.

The Equity Concern

Perhaps the most troubling limitation is one of access. A January 2026 study in the *American Journal of Managed Care* found that among 2,784 US hospitals using Epic, 62.6 percent had adopted ambient AI—but adoption was starkly uneven. Nonprofit hospitals adopted at 70.2 percent, while for-profit hospitals adopted at only 28.8 percent. Metropolitan hospitals led rural ones [30]. If ambient AI truly restores the physician-patient relationship and reduces errors linked to burnout, then the communities least likely to have access to safety-net hospitals, rural clinics, and under-resourced systems are precisely those that need it most. The technology risks widening the very disparities it could help close (see Figure 2.3).

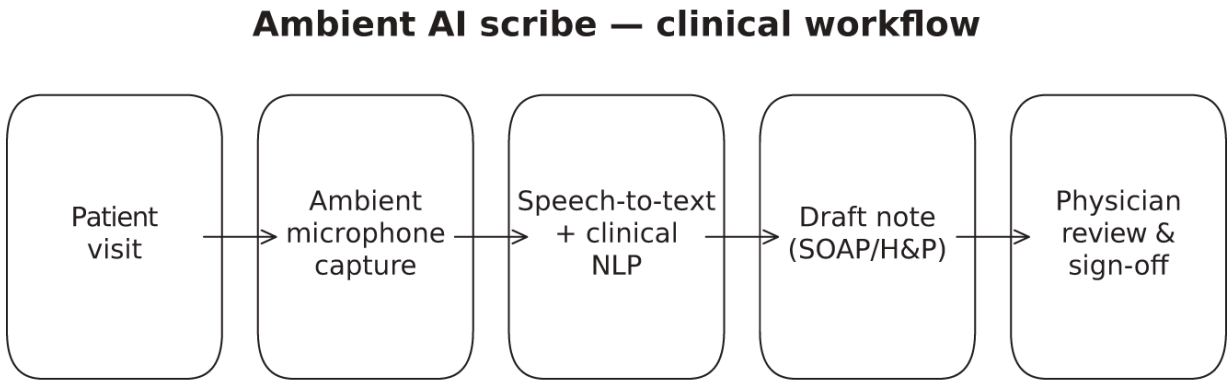


Figure 2.3 Ambient AI documentation restoring physician–patient connection

Regulatory and Legal Landscape

The regulatory landscape for ambient AI documentation exists in a peculiar gray zone. Most ambient AI scribes are not classified as medical devices by the Food and Drug Administration because they do not make clinical decisions; they generate draft text that a physician must review and sign. This distinction places them outside the FDA's traditional premarket review framework, unlike diagnostic AI tools such as diabetic retinopathy screeners or stroke detection algorithms, which do require FDA clearance. The practical consequence is that ambient AI has been deployed at enterprise scale across hundreds of health systems without the kind of regulatory scrutiny that accompanies a new drug, a new imaging device, or even a new clinical decision-support algorithm.

This regulatory void is not inherently problematic; in its current form, ambient AI is a productivity tool rather than a clinical tool—but it does raise questions as these systems evolve. Dragon Copilot, the successor to DAX Copilot, launched by Microsoft in March 2025, now includes capabilities beyond note generation: conversational orders, referral letters, and after-visit summaries [31]. As ambient AI begins to interact with clinical workflows, suggesting orders, summarizing patient histories for clinical decisions, or generating discharge instructions, the line between "productivity tool" and "clinical decision support" will blur, and the regulatory framework will need to adapt. The JAMIA 2025 survey of 43 health systems found that while 100 percent reported activity in ambient documentation, the most frequently cited barriers to broader AI adoption included immature tools, financial concerns, and regulatory uncertainty, the last of which directly reflects the absence of clear FDA guidance on this evolving category of technology [39].

HIPAA and Patient Consent

Every major ambient AI deployment requires that the physician obtain patient consent before initiating a recording. At the national level, HIPAA governs the handling of protected health information, and ambient AI platforms must comply with its security and privacy requirements. Audio recordings are generally not retained after the AI generates its note draft. This design choice reduces data risk but also limits the ability to audit AI

performance after the fact. State consent laws add complexity: in “two-party consent” states, both the patient and the physician must agree to the recording, whereas in “one-party consent” states, only one participant’s agreement is required. In practice, all major health systems have adopted explicit patient consent protocols regardless of state law, and patient declination rates have been consistently low, ranging from 6.4 percent to 7.2 percent in the UCLA RCT and as low as 0.014 percent at Intermountain Health [20,32].

Malpractice Implications

The malpractice question is both pressing and unresolved. When an AI-generated note contains an error, a hallucinated medication, a misattributed symptom, or an omitted physical exam finding, and the physician signs the note without catching the error, liability almost certainly rests with the physician. The physician remains the legal author of the note. But as ambient AI becomes ubiquitous and the expectation of physician review becomes routine, courts will inevitably face questions about shared liability: what responsibility does the AI vendor bear for systematic patterns of error? What duty does the health system have to monitor AI note quality? What standard of review should a “reasonable physician” apply to an AI-generated note? These questions have not yet been tested in litigation, but they will be.

The Joint Commission and CHAI Framework

In September 2025, the Joint Commission and the Coalition for Health AI (CHAI) together released the Responsible Use of AI in Healthcare (RUAIH) framework—a nonbinding yet influential set of seven governance steps for all 22,000-plus Joint Commission–accredited organizations. The framework calls for formalized AI governance structures with board-level oversight, patient privacy and transparency requirements, ongoing quality monitoring with vendor validation, voluntary safety reporting for AI incidents, risk and bias assessment, and role-specific training and authorization protocols [33]. CHAI has also launched a Model Card Registry, analogous to a “nutrition label” for AI, that allows health systems to document an AI model’s development, capabilities, limitations, and bias mitigation measures. Major systems, including Kaiser Permanente,

Cleveland Clinic, Providence, Stanford, Duke, and Memorial Sloan Kettering, have endorsed the registry [33].

Ethical Implications

The ethical terrain of ambient AI is more complex than its proponents often acknowledge. The technology raises questions that extend beyond privacy into the domains of consent, equity, power, and the nature of the clinical relationship itself.

Consent and Transparency

When a physician tells a patient that an AI tool will be listening to their conversation and generating a clinical note, what does the patient actually understand? The consent models currently in use are based on simple verbal notification, in which the physician explains the tool and the patient agrees or declines. But true informed consent requires understanding of what the technology does with the conversation data, how long it is retained, who has access, and its limitations. Patients who are elderly, have limited English proficiency, are in acute distress, or face power imbalances with their clinician may agree out of deference rather than understanding. The 8 percent of Kaiser patients who reported discomfort with the technology—and the 26 percent who were neutral—suggest that a meaningful minority may not be fully comfortable even in an institution known for patient engagement [6].

The power dynamics of opting out deserve scrutiny. A patient who declines ambient AI is, in effect, asking their physician to revert to a less efficient workflow. In a busy clinic, that request—however politely received—may subtly alter the dynamics of the encounter. The pressure to conform to the default is significant, and the default, in an institution that has deployed ambient AI enterprise-wide, is to be recorded.

Efficiency vs. Well-Being

The ECRI Institute, one of the nation’s leading patient safety organizations, has flagged a concern that bears directly on how health systems deploy ambient AI: the risk that efficiency gains are used to increase physician

productivity rather than improve physician well-being. Providence's approach—explicitly not requiring physicians to see more patients as a condition of using DAX Copilot—represents an intentional choice to direct the time savings toward clinician well-being rather than throughput [34]. But not every system has made that commitment. When the Providence RCT found that physicians naturally saw an average of 1.7 additional patients per day, it was presented as a voluntary, organic effect of reduced backlog—not a mandate [21]. The question is whether that distinction will hold as health systems face mounting financial pressures. If ambient AI becomes a mechanism for extracting more clinical work from the same number of physicians rather than restoring their humanity, it will have failed its deepest promise.

Equity and Access

The adoption data make clear that ambient AI is reaching well-resourced nonprofit and academic medical centers first, and the gap is already measurable. The patients most likely to benefit—those in under-resourced communities where physician burnout, understaffing, and access barriers are most severe—are the least likely to encounter the technology. This is not a hypothetical concern; it is the historical pattern of every healthcare technology innovation, from EHRs to telemedicine. Without deliberate policy intervention—including subsidies, technical assistance, and deployment incentives for safety-net systems—ambient AI will widen the care gap. The AJMC study showing that an AI-integrated readmission reduction program at Zuckerberg San Francisco General Hospital eliminated the racial equity gap in heart failure readmissions demonstrates what is possible when AI is deployed with equity at the center [35]. The question is whether the field will scale that model or default to the historical pattern of innovation flowing to the already privileged.

The Productivity Trap

Finally, there is a subtle but important risk embedded in the economics of ambient AI. If the technology saves a physician 30 minutes per day, that time can be redirected toward well-being—leaving early, mentoring a student, simply taking a breath between patients—or toward seeing more patients. The Providence RCT observed an average of 1.7 additional

patients per day as an organic workflow effect, not a mandate [21]. At Hattiesburg Clinic, 43 percent of physicians using a single vendor's tool reported they could add a patient if they wanted to [37]. Health system leaders under financial pressure may be tempted to capture this capacity as a productivity gain rather than return it to clinicians as restored well-being. If that happens—if ambient AI becomes a mechanism for squeezing more clinical hours from the same burned-out workforce—the technology will have failed the very people it was designed to serve.

Leadership Implementation Framework

For hospital and health system leaders considering ambient AI deployment—the CEOs, CIOs, and CMOs whose decisions will shape how this technology reaches clinicians—the research offers both a compelling case and hard-won implementation lessons. What follows is a practical framework drawn from the institutions that have done this at scale.

Start with the Why—and Make It About Clinicians

Every successful deployment began with a clear articulation that the primary goal was clinician well-being, not productivity. Providence's CMIO, Dr. Scott Smitherman, was explicit: "We did not fund the technology by requiring higher patient volumes…. Whether you're asking [clinicians] to pay the fee or whether you're asking them to virtually pay for it by seeing more patients…we did neither" [34]. That message mattered. When clinicians believe the technology is being deployed for their benefit—rather than as a mechanism to extract more work—adoption follows. Providence reports that approximately half of its primary care physicians now use DAX for the majority of their visits, and the system has expanded to emergency departments and inpatient settings [34].

Pilot Rigorously, Then Scale Deliberately

Kaiser Permanente's approach is instructive. Before enterprise deployment, TPMG conducted a 10-week pilot with more than 1,000 clinicians across 63,000 encounters, scoring AI-generated transcripts on a standardized quality rubric. Only after the quality assurance process confirmed acceptable accuracy did the system proceed to enterprise rollout [25].

Cleveland Clinic took a different but equally rigorous approach: a full year of evaluation across five ambient AI tools, involving approximately 250 physicians in more than 80 specialties, before selecting Ambience Healthcare for enterprise deployment. Within 15 weeks of rollout, more than 4,000 of 6,000 eligible physicians were active users—a 67 percent adoption rate that reflected the credibility built during the evaluation period [36].

Invest in Physician Champions

The institutions with the highest adoption rates universally identified physician champions as a critical success factor. At Hattiesburg Clinic, a 300-physician group practice in Mississippi, CEO Bryan Batson, MD, personally led the pilot. A pediatrician's unsolicited social media post thanking patients for their support of the new technology created organic demand that spread through the clinic faster than any institutional communication campaign could have [37]. At Stanford, Tait Shanafelt himself—the researcher who defined the burnout crisis—co-authored the deployment studies, lending both credibility and urgency to the effort [38]. The JMIR 2024 systematic review of AI implementation in hospitals identified "innovation champions" as a key enabler across all phases of deployment [39].

Training Must Be Clinical, Not Just Technical

The most effective training programs were peer-led and clinically embedded. Hattiesburg used "at-the-elbow" training by personnel with clinical AI experience—not by vendor representatives reading scripts [37]. UW Health employed implementation scientists and human factors experts through rapid Plan-Do-Study-Act (PDSA) cycles, a standard quality-improvement framework for testing changes in small, iterative rounds, contributing to the highest utilization rate (71 percent) of any ambient AI RCT to date [7]. Stanford used multimodal training—knowledge base articles, group sessions, individual sessions, and weekly ad hoc support—with the pilot yielding a 98 percent "easy to use" rating from physicians [38].

Governance Is Not Optional

The Joint Commission/CHAI framework provides a starting template, but leading institutions have gone further. Kaiser Permanente participates in the CHAI Model Card Registry. Cleveland Clinic's chief digital officer oversees AI strategy with the ACMIO and leads clinical evaluation. HCA Healthcare uses multidisciplinary governance boards that include operations, finance, supply chain, and clinical leaders—not just IT [40]. The VA has established a chief AI officer within its Office of Information and Technology and a National AI Institute within the Veterans Health Administration, operating through a hub-and-spoke governance model across its regional networks [41]. Every health system deploying AI in clinical settings needs a governance committee with board-level visibility, defined escalation pathways, and ongoing performance monitoring—not a one-time approval process.

Monitor, Measure, and Report Transparently

The most credible deployments published their results in peer-reviewed journals—including negative or ambiguous findings. The UCLA RCT reported that DAX did not achieve statistical significance on the primary outcome, even as it showed burnout benefits on secondary measures [20]. The Sutter Health study reported that improvements in burnout did not reach statistical significance [29]. This transparency builds trust—among clinicians, among patients, and among the broader healthcare community. TPMG maintains ongoing tracking of EHR time metrics, after-hours time, patient opt-out rates, and physician survey feedback. UW Health developed a live data dashboard for post-deployment monitoring of ambient AI performance across its 800 users. The VA requires AI use-case registration and an annual public inventory [41]. These monitoring practices should be considered minimum requirements, not aspirational goals.

Future Outlook

The ambient AI revolution in documentation is, in many respects, the opening act. The technology that listens to a physician-patient conversation and generates a note is proving its value—but the larger transformation will come from the convergence of ambient AI with diagnostic AI, predictive

AI, and care coordination AI into what might be called “unified clinical intelligence.”

The Diagnostic Partnership

The evidence for AI as a diagnostic partner is already striking. In a 2025 study published in Nature, Google’s AMIE system—a large language model optimized for clinical reasoning—achieved a top-10 diagnostic accuracy of 59.1 percent on challenging real-world cases from the NEJM Clinicopathological Conference, compared with 33.6 percent for unassisted clinicians. When physicians used AMIE as a thinking partner, they added the correct diagnosis to their list in 73 cases where they would otherwise have missed it [42]. For patients with rare diseases, the implications are profound: a JAMA Network Open study found that large language models correctly identified diagnoses in 13.3 percent of cases from the NIH’s Undiagnosed Diseases Network—more than double the historical 5.6 percent clinical review rate—for 3 cents per case [43].

Cancer Detection

AI mammography is moving from promising to proven. The ASSURE trial, published in *Nature Health* in 2025, analyzed more than 579,000 women across 109 community imaging sites and found that AI-supported screening detected 21.6 percent more cancers than standard mammography—with no racial or breast-density-based disparities in detection rates [44]. The Swedish MASAI trial showed approximately 30 percent more cancers detected, with radiologist workload reduced by nearly half [45]. In China, the PANDA AI tool reviewed more than 180,000 CT scans and identified approximately 24 pancreatic cancers, 14 at an early stage—including the case of Mr. Qiu, a 57-year-old retired brick worker whose tumor was flagged by AI and successfully removed. “You can confidently say AI has saved their lives,” said Dr. Zhu Kelei, the treating surgeon [46].

What 2030 Might Look Like

Imagine a physician’s day in 2030. She walks into the exam room and greets her patient, with no laptop open and no keyboard between them. The ambient AI captures their conversation and generates not just a note but also a preliminary differential diagnosis informed by the patient’s full history,

recent labs, imaging trends, and population-level risk factors. An AI-driven care coordination system has already flagged that this patient missed a follow-up colonoscopy and has identified a transportation barrier; a social work referral is queued. After the visit, the physician reviews the AI-generated note, which has been checked against an internal evidence base for medication interactions and guideline adherence, makes three edits, and signs it. By 5:15 p.m., her charts are done. She picks up her children from soccer practice.

That future is not science fiction. Every component of that scenario exists today in some form, deployed at one or more major health systems. The challenge over the next five years is to integrate the ambient documentation layer, the diagnostic reasoning layer, the care coordination layer, and the predictive analytics layer into a seamless clinical intelligence platform.

The Standard-of-Care Question

There is a final question that health system leaders, malpractice attorneys, and regulators will eventually need to answer: will *not* having ambient AI become a liability? If the evidence continues to show that ambient AI reduces burnout, reduces medical errors linked to burnout, and improves documentation quality, at what point does deploying it become the expected standard and failing to deploy it a deviation from accepted practice? We are not there yet. But the trajectory of the evidence suggests that the question is not *if* but *when*. The speed of adoption underscores the urgency: the JAMIA survey of 43 health systems found that 100 percent reported ambient AI activity, and 62.6 percent of Epic-using U.S. hospitals had adopted ambient AI by mid-2025, a pace of adoption that rivals the EHR incentive surge under HITECH a decade ago [30,39]. When a technology moves from pilot to near-universal adoption in less than three years, the standard-of-care conversation is not far behind.

Key Takeaways

Burnout remains at crisis levels. Despite modest improvements, 43–49 percent of U.S. physicians report burnout symptoms, and physicians are 82.3 percent more likely to experience burnout than

other U.S. workers. Documentation burden is the single most consistent driver.

Ambient AI has the strongest evidence base of any burnout intervention. Three randomized controlled trials and multiple large-scale studies demonstrate that ambient AI reduces burnout scores, documentation time, after-hours charting, and cognitive load, with the UW Health trial showing a number needed to treat of just 1.68.

Real-world deployment at scale is feasible and impactful. Kaiser Permanente's deployment across 7,260 physicians and 2.5 million encounters saved an estimated 15,791 hours of documentation time, with 88 percent of physicians reporting positive effects on visit interactions and 47 percent of patients saying their doctor spent less time on the computer.

Note quality requires ongoing vigilance. AI-generated notes contain hallucinations at measurable rates (up to 31 percent in one study), and the risk of automation complacency, when physicians review AI notes with decreasing scrutiny, is real and must be actively managed.

AI alone cannot fix burnout. The documentation burden is one driver among many. Staffing, compensation, autonomy, and organizational culture must all be addressed. Systems that use ambient AI to increase throughput rather than restore well-being will undermine the technology's promise.

Equity must be built in, not bolted on. Ambient AI adoption is concentrated in well-resourced nonprofit and metropolitan hospitals. Without deliberate policy action, the technology will widen existing disparities in care quality and clinician well-being.

The convergence is coming. Ambient documentation is the first wave. The integration of documentation AI, diagnostic AI, predictive AI, and care coordination AI into unified clinical intelligence platforms will define the next era of medicine and raise the standard of care for every health system.

References

1. Sinsky C, Colligan L, Li L, Prgomet M, Reynolds S, Goeders L, et al. Allocation of physician time in ambulatory practice: a time and motion study in 4 specialties. Ann Intern Med. 2016. Available from: https://adfm.org/media/1476/ann-2016-time-study.pdf.

2. Arndt BG, Beasley JW, Watkinson MD, Temte JL, Tuan WJ, Sinsky CA, et al. Tethered to the EHR: primary care physician workload assessment using EHR event log data and time-motion observations. Ann Fam Med. 2017; 15(5):419-426. doi:10.1370/afm.2121.

3. Tawfik DS, Profit J, Morgenthaler TI, Satele DV, Sinsky CA, Dyrbye LN, et al. Physician burnout, well-being, and work unit safety grades in relationship to reported medical errors. Mayo Clin Proc. 2018; 93(11):1571-1580. doi:10.1016/j.mayocp.2018.05.014. PMID 30001832. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC6258067.

4. White T. Medical errors may stem more from physician burnout than unsafe medical workplace conditions. Stanford Medicine News Center. 2018 Jul 9 [cited 2026 Jul 4]. Available from: https://med.stanford.edu/news/all-news/2018/07/medical-errors-may-stem-more-from-physician-burnout.html.

5. Physician Burnout and Suicide: A Call to Action. Oncology Practice Management. 2023 Sep; 13(9) [cited 2026 Jul 4]. Available from: https://oncpracticemanagement.com/issues/2023/september-2023-vol-13-no-9/3253:physician-burnout-and-suicide-a-call-to-action.

6. Tierney AA, Gayre G, Hoberman B, Mattern B, Ballesca M, Wilson Hannay SB, et al. Ambient artificial intelligence scribes: learnings after 1 year and over 2.5 million uses. NEJM Catalyst. 2025 Mar 31. doi:10.1056/CAT.25.0040.

7. Afshar M, Ryan Baumann M, Resnik F, Hintzke J, Gravel Sullivan A, Wills G, et al. A pragmatic randomized controlled trial of ambient

artificial intelligence to improve health practitioner well-being. NEJM AI. 2025 Nov 26; 2(12). doi:10.1056/aioa2500945.

8. Shanafelt TD, West CP, Sinsky C, Trockel M, Tutty M, Wang H, et al. Changes in burnout and satisfaction with work-life integration in physicians and the general US working population between 2011 and 2023. Mayo Clin Proc. 2025 Jul; 100(7):1142-1158. doi:10.1016/j.mayocp.2024.11.031.

9. Advisory Board Daily Briefing. Physician burnout. 2024 Jan 31 [cited 2026 Jul 4]. Available from: https://www.advisory.com/daily-briefing/2024/01/31/physician-burnout.

10. Medscape. Despite positive shift unveiled in annual physician burnout and depression survey, there's still a lot of work to do. PR Newswire. 2024 Jan 24 [cited 2026 Jul 4]. Available from: https://www.prnewswire.com/news-releases/new-medscape-report-reveals-progress-among-physician-burnout-depression-302043454.html.

11. American Medical Association. AMA Organizational Biopsy® 2024 national physician comparison report. 2025 May. Available from: https://www.ama-assn.org/practice-management/physician-health/us-physician-burnout-hits-lowest-rate-covid-19.

12. Physicians Foundation. 2025 well-being survey: the state of America's physicians [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://physiciansfoundation.org/research/the-state-of-americas-physicians-2025-wellbeing-survey.

13. Ligibel JA, Goularte N, Berliner JI, Bird SB, Brazeau CMLR, Rowe SG, et al. Well-being parameters and intention to leave current institution among academic physicians. JAMA Netw Open. 2023 Dec 1; 6(12):e2347894. doi:10.1001/jamanetworkopen.2023.47894.

14. Association of American Medical Colleges. The complexities of physician supply and demand: projections from 2021 to 2036. 2024 Mar. Available from: https://www.aamc.org/news/press-releases/new-aamc-report-shows-continuing-projected-physician-shortage.

15. Han S, Shanafelt TD, Sinsky CA, Awad KM, Dyrbye LN, Fiscus LC, et al. Estimating the attributable cost of physician burnout in the United States. Ann Intern Med. 2019; 170(11):784-790. doi:10.7326/M18-1422.

16. Holmgren AJ, Sinsky CA, Rotenstein L, Apathy NC. National comparison of ambulatory physician electronic health record use across specialties. J Gen Intern Med. 2024 Nov; 39(14):2868-2870. doi:10.1007/s11606-024-08930-4.

17. Arndt BG, Micek MA, Rule A, Shafer CM, Baltus JJ, Sinsky CA. More tethered to the EHR: EHR workload trends among academic primary care physicians, 2019-2023. Ann Fam Med. 2024; 22(1):12-18. doi:10.1370/afm.3047.

18. Barr W, Peterson L, Fleischer S. Pajama time: the association of EHR documentation time with family medicine resident outcomes. Ann Fam Med. 2024; 22(Suppl 1):6628. doi:10.1370/afm.22.s1.6628. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC11627813.

19. Li LZ, Yang P, Singer SJ, Pfeffer J, Mathur MB, Shanafelt T. Nurse burnout and patient safety, satisfaction, and quality of care: systematic review and meta-analysis. JAMA Netw Open. 2024; 7(11):e2443059. doi:10.1001/jamanetworkopen.2024.43059.

20. Lukac PJ, Turner W, Vangala S, Chin AT, Khalili J, Shih YT, et al. Ambient AI scribes in clinical practice: a randomized trial. NEJM AI. 2025 Nov 26; 2(12). doi:10.1056/aioa2501000.

21. Wendt SJ, Dinh CT, Sutcliffe M, Jones K, Scanlan JM, Smitherman JS. Deploying ambient clinical intelligence to improve care: a research article assessing the impact of Nuance DAX on documentation burden and burnout. Future Healthc J. 2025; 12(2):100450. doi:10.1016/j.fhj.2025.100450. PMID 40896303.

22. Olson KD, Meeker D, Troup M, Barker TD, Nguyen VH, Manders JB, et al. Use of ambient AI scribes to reduce administrative burden and professional burnout. JAMA Netw Open. 2025 Oct 1; 8(10):e2534976. doi:10.1001/jamanetworkopen.2025.34976.

23. Zarefsky M. Ambient scribes and AI feel like magic in health care. American Medical Association [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/ambient-scribes-and-ai-feel-magic-health-care.

24. Bruce G. Make me a better doctor — Kaiser Permanente's AI playbook. Becker's Hospital Review [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.beckershospitalreview.com/healthcare-information-technology/ai/make-me-a-better-doctor-kaiser-permanentes-ai-playbook.

25. Cain CH, Davis AC, Broder B, Chu E, DeHaven AH, Domenigoni A, et al. Quality assurance during the rapid implementation of an AI-assisted clinical documentation support tool. NEJM AI. 2025 Mar 27. doi:10.1056/AIcs2400977.

26. Feldheim B. AI scribes save 15,000 hours and restore the human side of medicine. American Medical Association [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/ai-scribes-save-15000-hours-and-restore-human-side-medicine.

27. Palm E, Manikantan A, Mahal H, Belwadi SS, Pepin ME. Assessing the quality of AI-generated clinical notes. Front Artif Intell. 2025; 8:1691499. doi:10.3389/frai.2025.1691499.

28. Andrews M. Doctors increasingly see AI scribes in a positive light. But hiccups persist. KFF Health News [Internet]. 2026 Jan 27 [cited 2026 Jul 4]. Available from: https://kffhealthnews.org/news/article/ambient-ai-scribes-doctor-appointments-note-taking-ehr-epic.

29. Stults CD, Deng S, Martinez MC, Wilcox J, Szwerinski N, Chen KH, et al. Evaluation of an ambient artificial intelligence documentation platform for clinicians. JAMA Netw Open. 2025 May 1; 8(5):e258614. doi:10.1001/jamanetworkopen.2025.8614. PMID 40314951.

30. Yang F, Graetz I. Ambient AI tool adoption in US hospitals and associated factors. Am J Manag Care. 2026 Jan 27; 32(1).

31. Microsoft. Dragon Copilot launch. 2025 Mar 3. Available from: https://news.microsoft.com/source/2025/03/03/microsoft-dragon-copilot.

32. Ma SP, Liang AS, Shah SJ, Smith M, Jeong Y, Devon-Sand A, et al. Ambient artificial intelligence scribes: utilization and impact on documentation time. J Am Med Inform Assoc. 2024 Dec 17; 32(2):381-385. doi:10.1093/jamia/ocae304.

33. The Joint Commission; Coalition for Health AI. Responsible use of AI in healthcare (RUAIH) framework. 2025 Sep 17.

34. Providence Office. DAX Copilot: “Life-changing” technology. Providence Blog [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://blog.providence.org/blog/dax-copilot-life-changing-technology.

35. Reducing readmissions in the safety net through AI and automation. Am J Manag Care. 2025; 31(3):142-148. doi:10.37765/ajmc.2025.89697.

36. Cleveland Clinic. Less typing, more talking: how ambient AI is reshaping clinical workflow. Consult QD [Internet]. 2025 [cited 2026 Jul 4].

37. American Medical Association. Hattiesburg clinic doctors say ambient AI lowers stress, burnout. AMA [Internet]. 2024 [cited 2026 Jul 4].

38. Shah SJ, Devon-Sand A, Ma SP, Jeong Y, Crowell T, Smith M, et al. Ambient artificial intelligence scribes: physician burnout and perspectives on usability and documentation burden. J Am Med Inform Assoc. 2025 Feb 1; 32(2):375-380. doi:10.1093/jamia/ocae295.

39. Poon EG, Lemak CH, Rojas JC, Guptill J, Classen D. Adoption of artificial intelligence in healthcare: survey of health system priorities, successes, and challenges. J Am Med Inform Assoc. 2025 Jul; 32(7):1093-1100. doi:10.1093/jamia/ocaf065.

40. American Hospital Association. Smarter, safer hospitals: how HCA Healthcare is using AI to redefine patient safety. AHA [Internet]. 2025 [cited 2026 Jul 4].

41. US Department of Veterans Affairs. Building the future: VA's strategy for adopting high-impact artificial intelligence. VA [Internet]. 2026 [cited 2026 Jul 4].

42. McDuff D, Schaekermann M, Tu T, Palepu A, Wang A, Garrison J, et al. Towards accurate differential diagnosis with large language models. Nature. 2025 Jun; 642(8067):451-457. doi:10.1038/s41586-025-08869-4.

43. Shyr C, Cassini TA, Tinker RJ, Byram KW, Embí PJ, Bastarache L, et al. Large language models for rare disease diagnosis at the Undiagnosed Diseases Network. JAMA Netw Open. 2025. doi:10.1001/jamanetworkopen.2025.28538.

44. Louis LD, Wakelin EA, McCabe MP, Ng AY, Kim JG, Lee CI, et al. Equitable impact of an AI-driven breast cancer screening workflow in real-world US-wide deployment. Nat Health. 2025 Nov 17. doi:10.1038/s44360-025-00001-0.

45. Gommers J, Hernström V, Josefsson V, Sartor H, Schmidt D, Hjelmgren A, et al. Interval cancer, sensitivity, and specificity comparing AI-supported mammography screening with standard double reading without AI in the MASAI study: a randomised, controlled, non-inferiority, single-blinded trial. Lancet. 2026 Feb 7; 407(10527):505-514. doi:10.1016/S0140-6736(25)02464-X.

46. New York Times. In China, AI finds pancreatic cancer that doctors may miss. 2026 Jan 2.

47. American Nurses Foundation. 2023 pulse on the nation's nurses survey. 2023 May.

48. National Council of State Boards of Nursing. 2024 national nursing workforce study. 2025 Apr 17.

49. NSI Nursing Solutions. 2025 national health care retention & RN staffing report.

50. KLAS Research. Clinician turnover 2024. KLAS [Internet]. 2024 [cited 2026 Jul 4].

51. Commonwealth Fund. A poor prognosis? More than one-third of burned-out U.S. primary care physicians plan to stop seeing patients. Commonwealth Fund Blog [Internet]. 2024 Dec 6 [cited 2026 Jul 4]. Available from: https://www.commonwealthfund.org/blog/2024/poor-prognosis-more-one-third-burned-out-us-primary-care-physicians-plan-stop-seeing.

52. Kaiser Permanente. Kaiser Permanente improves member experience with AI-enabled clinical technology. about.kaiserpermanente.org [Internet]. 2024 Aug 14 [cited 2026 Jul 4]. Available from: https://about.kaiserpermanente.org/news/press-release-archive/kaiser-permanente-improves-member-experience-with-ai-enabled-clinical-technology.

53. Shah SJ, Crowell T, Jeong Y, Devon-Sand A, Smith M, Yang B, et al. Physician perspectives on ambient AI scribes. JAMA Netw Open. 2025 Mar 3; 8(3):e251904. doi:10.1001/jamanetworkopen.2025.1904.

54. You JG, Dbouk RH, Landman A, Ting DY, Dutta S, Wang JC, et al. Ambient documentation technology in clinician experience of documentation burden and burnout. JAMA Netw Open. 2025 Aug 1; 8(8):e2528056. doi:10.1001/jamanetworkopen.2025.28056.

55. Pearlman K, Wan W, Shah SD, Laiteerapong N. Use of an AI scribe and electronic health record efficiency. JAMA Netw Open. 2025 Oct.

56. Duggan MJ, Gervase J, Schoenbaum A, Hanson W, Howell JT 3rd, Sheinberg M, Johnson KB. Clinician experiences with ambient scribe technology. JAMA Netw Open. 2025 Feb 3; 8(2):e2460637. doi:10.1001/jamanetworkopen.2024.60637.

AI in Emergency Medicine: What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution, First Edition. Harvey Castro.
© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

3
Data, Bias, and the Human Cost of Errors

My mother arrived in the United States from Colombia at 13 years old. She became a mother at 16. I was raised on food stamps, government cheese, and powdered milk that came in boxes with no brand name. By the time I was 4 years old, I was a latchkey kid, with a key on a string around my neck, letting myself into an empty apartment after school because my mother was working her second or third job of the day.

We didn’t go to the doctor. That was the rule, and it wasn’t about toughness; it was about arithmetic. A doctor’s visit cost money we didn’t have. So my mother taught me home remedies: herbs, rest, prayer, and the stubborn hope that whatever was wrong would pass on its own. Unless I was actively dying, the hospital was not an option.

When she did seek care, for the things that couldn’t be prayed away, the language barrier turned every visit into a guessing game. English was her second language, medical English was nobody’s first, and no one in those clinics spoke Spanish. She would nod at explanations she didn’t fully understand, take prescriptions she couldn’t read, and leave with a vague sense that something had been decided about her body by people who never asked what she thought.

I didn’t have a name for this as a child. I just knew that the healthcare system felt like it was designed for someone else.

Decades later, I sat in my office reading Obermeyer et al. in *Science* [1], the landmark 2019 study demonstrating that a widely used commercial algorithm systematically under-referred Black patients for care management programs. The algorithm used healthcare spending as a proxy for health need. But spending reflects access, not illness. People who can’t afford to go to the doctor don’t generate spending data. The algorithm learned, accurately, that these patients cost less. It concluded, catastrophically, that they were healthier.

I didn’t need the paper to explain the mechanism. I had lived it. My mother didn’t generate healthcare spending data either. She worked three jobs and avoided the clinic. By the logic of that algorithm, she was perfectly healthy. She was an overworked immigrant with no safety net who would eventually discover, by scanning herself with an ultrasound machine she operated for other people, that she had two cancers growing inside her.

The bias in AI doesn’t come from malice. It comes from data, and data reflects the system that generated it. When 15 percent of AI training data represents the low-income population, a figure cited in the research this chapter examines, the model isn’t just incomplete. It is encoding the very invisibility my mother experienced into its architecture. The algorithm doesn’t see you for the same reason the clinic didn’t see her: you were never in the dataset because you were never in the door.

This chapter is about what happens when we build powerful tools on broken foundations. I write it not as an abstract policy concern but as the son of a woman who the data would have missed entirely.

Executive Summary

Every algorithm begins as a mirror. It reflects the world that built it: the hospitals whose records filled its training sets, the patients who happened to seek care, the billing codes that named their suffering. The trouble is that the world mirrored in healthcare data is not a fair one. It carries decades of structural inequality: underfunded clinics that generated thinner records, patients who avoided hospitals they couldn’t afford, and measurement tools that were never calibrated for bodies that looked different from those in the founding studies. When artificial intelligence learns from that world, it doesn’t simply inherit these distortions. It scales them.

Chapter 3 confronts the most urgent vulnerability in the AI healthcare revolution: the systemic ways in which flawed data and biased algorithms cause real, measurable harm to real patients. The evidence is not theoretical. A landmark study published in *Science* in 2019 found that a single commercial algorithm (the Optum Impact Pro tool, used to guide care decisions for more than 70 million Americans, part of a broader class of similar tools affecting roughly 200 million people nationally) was enrolling

Black patients in supplemental care programs at less than half the rate their actual health burden warranted [1]. The same pattern runs through dermatology, cardiology, pulmonology, oncology, and emergency medicine: tools that work better for some patients than others, with no transparency, no accountability, and often no one knowing.

This chapter traces the problem from its roots in training data to its consequences in patient outcomes and from individual case studies to systemic regulatory failures. It examines three real-world cases: IBM Watson for Oncology's unsafe cancer treatment recommendations; the Epic Sepsis Model's near-collapse in external validation; and a surgical navigation system that contributed to patient strokes after an AI module was added primarily as a marketing tool. It also confronts the cognitive trap of automation bias: a biased AI model can reduce clinician diagnostic accuracy by more than 11 percentage points, even when explanations are provided [2].

The chapter also charts the path forward. Recalibration of the Optum algorithm eliminated the racial enrollment gap (see Figure 3.1). Fine-tuning dermatology AI on diverse images eliminated performance disparities. The eGFR race correction, embedded for decades, has been removed from clinical calculators. What the evidence shows, unmistakably, is that the problem is not technical intractability. It is institutional will. For healthcare leaders prepared to govern AI with the same rigor they apply to medication safety and credentialing, the tools, frameworks, and evidence already exist. The question is whether they choose to use them.

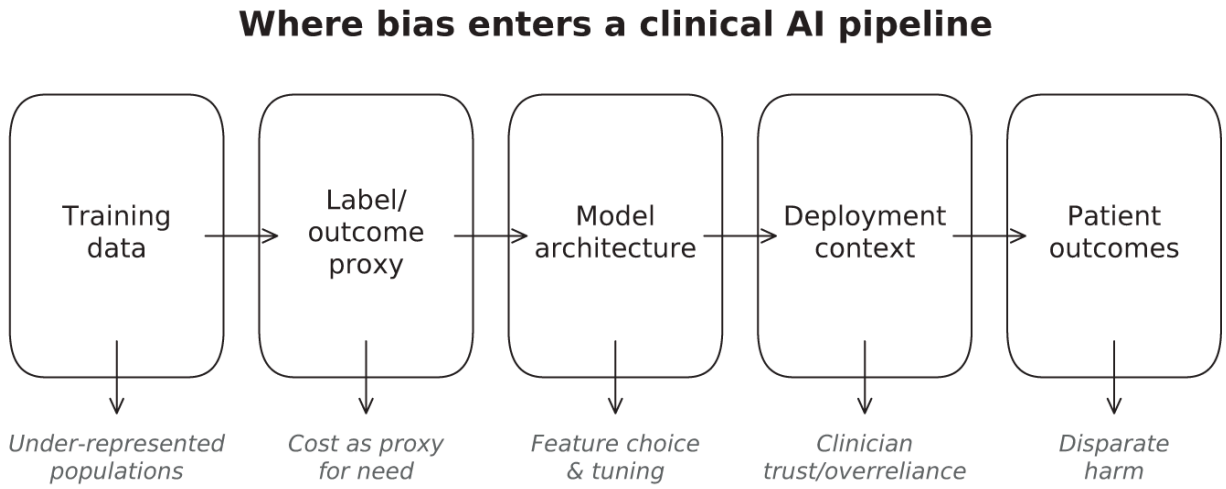


Figure 3.1 How biased data flows through the algorithmic pipeline to produce biased outcomes

Current Evidence

The scientific literature documents this harm in careful, methodical detail. Three case studies illustrate how data shapes the tools we build and how those tools can amplify the inequities embedded in the training data.

The Optum Algorithm: When Cost Data Becomes a Proxy for Health

The *Science* paper arrived in October 2019 with the kind of precision that makes policy arguments impossible to dodge. Ziad Obermeyer and colleagues at the University of California, Berkeley, along with collaborators from the University of Chicago, had spent months examining a commercial algorithm used by many of the nation's most prominent accountable-care organizations: a tool that was helping decide which patients received enhanced care management services [1]. The algorithm's logic was straightforward: predict who will cost the most in the coming year, because those patients probably need the most care. The problem was that this premise, apparently reasonable, was hiding an injustice.

At any given risk score, Black patients were measurably sicker than white patients. They carried more chronic conditions, had worse biomarkers, and showed more signs of uncontrolled disease. Yet they were being flagged for extra care at dramatically lower rates, because decades of structural inequity

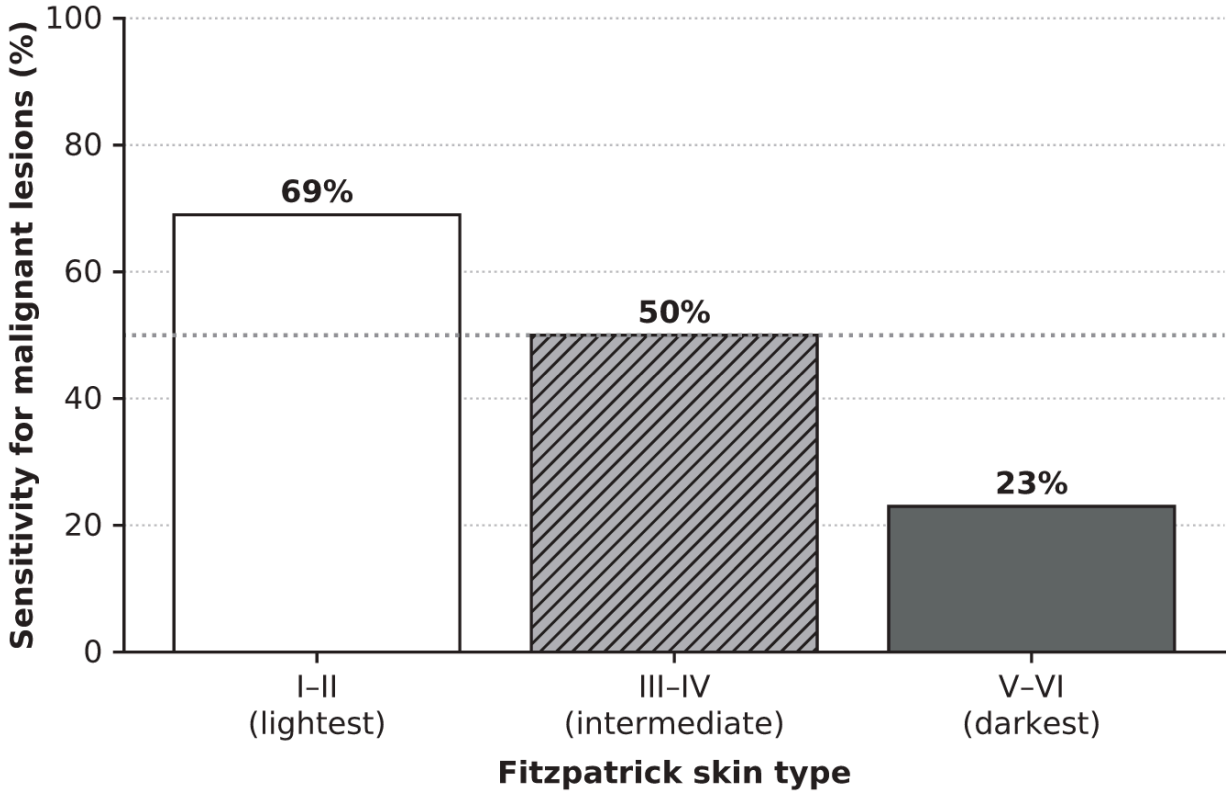
meant that equally sick Black patients had historically spent less on healthcare. This was not because they needed less but because they had less access to it. The algorithm learned from cost data, and cost data had encoded the entire history of American healthcare discrimination into a single number. The result: Black patients made up only 17.7 percent of those automatically enrolled in intensive care management programs. If the algorithm had used direct health indicators instead of costs, that figure would have risen to 46.5 percent, more than doubling Black enrollment [1].

The ACM companion paper noted this algorithm guided decisions for more than 70 million people nationally [1]. The scale is the point. A biased human clinician affects hundreds of patients over a career. A biased algorithm deployed across millions of clinical interactions amplifies that harm to a population level in every single computational cycle.

The dermatology evidence makes the mechanism painfully concrete. Skin cancer diagnosis depends on pattern recognition: the kind of visual learning at which AI has genuinely excelled. But pattern recognition requires patterns, and the patterns in dermatological training datasets have overwhelmingly come from patients with light skin tones. When Roxana Daneshjou and her colleagues at Stanford Medicine published their 2022 study in *Science Advances*, they had assembled the first publicly available, biopsy-confirmed image dataset with meaningful representation of all Fitzpatrick skin types: 656 images spanning light to dark complexions [3]. The results were stark. DeepDerm, one of the top-performing dermatology AI systems, showed a sensitivity for malignant lesions of 69 percent in patients with light skin types. For patients with the darkest skin tones, that sensitivity fell to 23 percent (see Figure 3.2) [3]. More than three-quarters of malignant lesions in darker-skinned patients went undetected. An AI tool that offered something approaching dermatologist-level screening to light-skinned patients was offering near-nothing to dark-skinned patients.

The structural cause was not mysterious. The International Skin Imaging Collaboration dataset, one of the most widely used dermatological image repositories in the world, draws predominantly from Europe and Australia, reflecting populations of primarily light skin. A systematic review found that across 136 dermatology AI studies, only 1.47 percent included participants from South America or Africa [4]. The Fitzpatrick 17k dataset (16,577 clinical images from dermatology atlases) confirmed the same

pattern: light skin types substantially outnumbered dark [5]. Models trained on these images don't generalize to the populations most likely to lack access to in-person dermatological care. Three billion people worldwide live without access to dermatologists. The patients most likely to encounter AI-based triage are often the ones for whom AI works worst [3].



Source: Daneshjou et al., Science Advances (2022).
Middle category interpolated; paper reports 69% (light) and 23% (darkest).

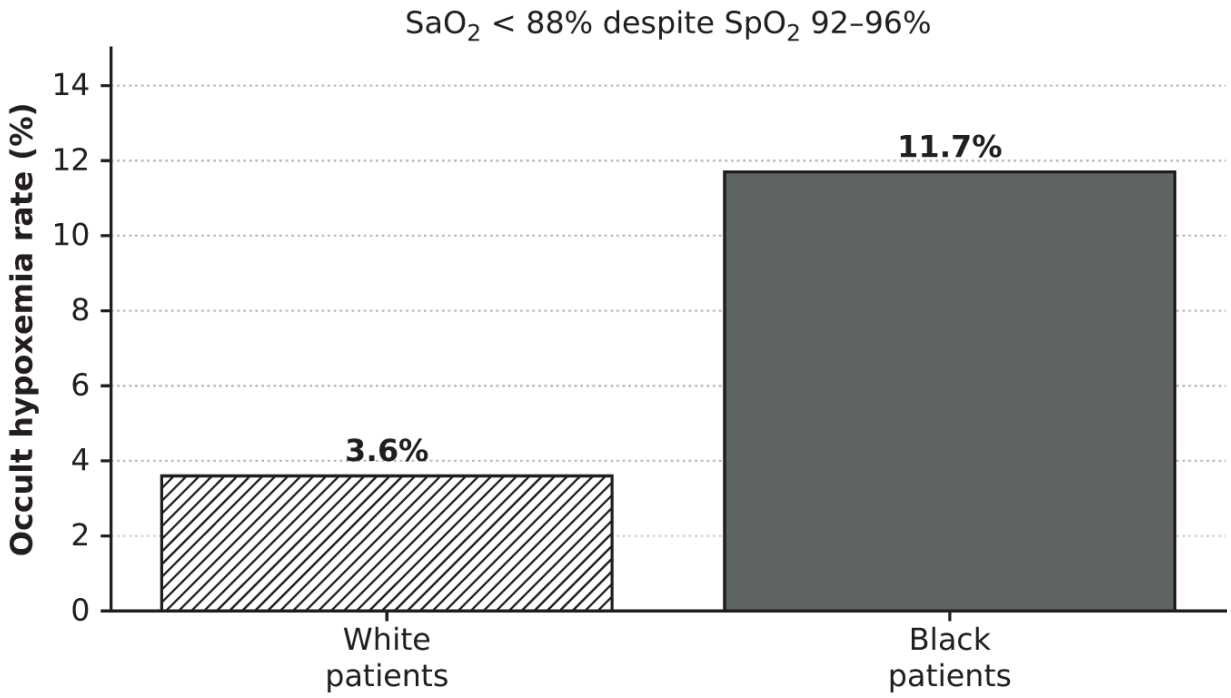
Figure 3.2 AI diagnostic accuracy declining across the Fitzpatrick skin tone scale

Pulse Oximetry: Decades of Documented Harm

Pulse oximetry tells a grimmer version of the same story, because here the bias predates AI entirely. Since at least 1990, and arguably as early as 1985, researchers had documented that pulse oximeters systematically overestimate oxygen levels in patients with darker skin pigmentation (see Figure 3.3) [6]. Melanin absorbs light similarly to oxygenated hemoglobin, producing a false reading. Despite this, documented in the peer-reviewed literature for more than three decade, no major manufacturer incorporated

corrected calibration algorithms. The FDA issued a safety communication in 2021.

When COVID-19 arrived, this 30-year-old uncorrected flaw became a treatment-access crisis. Fawzy and colleagues analyzed 7,126 COVID-19 patients at Johns Hopkins and published findings in *JAMA Internal Medicine* in 2022: Black patients had a 29 percent lower hazard of having their treatment eligibility, with a median delay of 7.0 hours compared to 5.3 hours for white patients [7]. The treatments at stake, remdesivir and dexamethasone, depended on pulse oximetry thresholds for access. Among the 451 patients whose eligibility was never recognized at all, 247 were Black (54.8 percent) and 122 were non-Black Hispanic (27.1 percent), while only 78 were white (17.3 percent) [7].



Source: Sjoding et al., NEJM (Dec 2020) — racial bias in pulse oximetry measurement.

Figure 3.3 Pulse oximeter racial bias: accurate readings on light skin, false readings on dark skin

Sudat and colleagues examined 43,753 concurrent arterial blood gas and pulse oximetry measurement pairs at Sutter Health in California and found that pulse oximetry overestimated oxygenation by 1 percent more in non-Hispanic Black patients than non-Hispanic white patients. That overestimation was associated with a 3.1 percentage-point lower probability

of hospital admission, a 3.1 percentage-point lower probability of dexamethasone treatment, and a staggering 278.5 additional minutes waiting for supplemental oxygen [8]. Wong and colleagues, in a study of 87,971 patients across five databases, confirmed that hidden hypoxemia in Black, Asian, and Hispanic patients was associated with mortality, organ dysfunction, and abnormal laboratory results [9]. A meta-analysis of 32 studies found a pooled mean bias of 1.52 percent overestimation in patients identified as Black or African American [10]. In critical care medicine, where oxygen saturation thresholds of 88, 92, or 94 percent determine whether a patient receives a potentially life-saving intervention, a systematic 1.52 percent overestimation is not a statistical footnote. It is a treatment barrier. In my experience, the patients most harmed by these gaps are precisely those with the least access to advocates who can navigate the system for them. When the algorithm fails silently, no one files a complaint. The data just quietly disappears.

Cardiovascular AI: The Gap in Gender Representation

Cardiovascular AI carries its own gender-stratified failures. Straw, Rees, and Nachev published findings in *Journal of Medical Internet Research* in 2023 showing that in 13 of 16 experiments across cardiac datasets, AI models produced significantly higher false negative rates for female patients, missing heart disease in women at rates that directly mirror the known clinical pattern of delayed cardiac diagnosis in women [11]. Only 3 of 60 peer-reviewed cardiac AI papers they examined had investigated sex differences in algorithm performance at all. None had investigated racial differences [11].

In 2025, Krieger and colleagues published findings in the *European Heart Journal* that should have shocked cardiology: when Generative Pre-trained Transformer 4 (GPT-4, an advanced AI language model developed by OpenAI) was prompted to generate cardiovascular case scenarios, 90 percent of the heart failure cases it generated were male, against a real-world prevalence of approximately 50 percent [12]. In diagnostic tasks using real patient notes from the MIMIC-IV database, GPT-4 showed significantly lower accuracy for female patients presenting with myocardial infarction and aortic dissection and more often misdiagnosed female MI patients with anxiety [12]. The large language model had absorbed

cardiology's historical gender blind spot and encoded it at inference-time scale.

Data Quality and Structural Bias

The data quality problems underlying these biases are structural and pervasive. A systematic review of EHR data quality in critical care settings found that missing data rates exceeded 80 percent for some variables in the MIMIC-III database and that more than 40 percent of the predictive features selected by machine learning algorithms were not clinical values at all but indicators of data absence [13]. The model was learning which patients had missing records, not which patients were sick. Copy-paste prevalence reached 82 percent in residents' progress notes and 74 percent in attending physicians' notes in the same review, systematically corrupting the natural language inputs that feed AI clinical documentation tools [13]. EHR-related medication errors comprised 34 percent of all ICU medication errors, with one-third having life-threatening potential [13].

Geographic concentration compounds the demographic problem. A Stanford University study cited in WHO regulatory guidance found that 71 percent of patient data from just three U.S. states trains most AI diagnostic tools used across the country [14]. Nearly half (44 percent) of published U.S. clinical AI studies did not report the ethnicity composition of their training datasets, according to a systematic review of approximately 400 published models [15]. A scoping review of 692 FDA-approved AI medical devices found that only 25 (3.6 percent) provided any information on the race or ethnicity of tested users, and 77.9 percent provided no demographic data whatsoever in their approval summaries [16].

These implications compound. AI systems often detect race even when it's never provided as an input. A *Lancet Digital Health* study found that deep learning models trained on chest X-rays, CT scans, and mammograms could predict patient self-reported race with AUC values of 0.91 to 0.99 across imaging modalities, performance sustained after image degradation, cropping, and corruption [17]. A 2025 study in *BMJ Health and Care Informatics* showed that vital signs alone, including heart rate, oxygen saturation, blood pressure, and respiratory rate, could predict race with an AUC of 0.74 between white and Black patients, with pulse oximetry identified as the most influential variable [18]. Any clinical decision support

system trained on EHR vital signs data carries an embedded racial signal, whether or not it was ever intended to. What I tell residents is that an algorithm's accuracy score is always the best-case number. It was measured on the patients who looked like the training set. The patients in your clinic may not.

Real-World Case Studies

These failures are not academic arguments. They represent real patients harmed by systems deployed before the evidence was clear.

IBM Watson for Oncology: The Overpromise

In 2015, IBM's marketing described Watson for Oncology as a system that would help defeat cancer. By 2018, internal documents told a different story. STAT News obtained slide decks from IBM Watson Health's deputy chief health officer presenting "multiple examples of unsafe and incorrect treatment recommendations" generated by the system [19].

The root failure was the training data. Watson had not learned from real patients. It had been trained primarily on synthetic, hypothetical cancer cases created by a small number of physicians and IBM engineers at Memorial Sloan Kettering Cancer Center: cases that reflected the clinical preferences of a handful of New York specialists, not evidence-based guidelines generalizable to the world. One internal document cited an example that captured the danger precisely: Watson recommended that a 65-year-old man with lung cancer and evidence of severe bleeding receive bevacizumab, a drug that carries an FDA black-box warning explicitly contraindicated in patients with serious bleeding events. Eric Topol of the Scripps Translational Science Institute, upon learning of the example, stated: "Any time you have an algorithm that makes a recommendation that's dangerous, that's extremely worrisome" [19].

Watson had been deployed at 230 hospitals in more than 50 countries before any independent, peer-reviewed clinical trial had examined its safety or efficacy. A collaboration with MD Anderson Cancer Center consumed four years and $62 million before being terminated without ever entering clinical use. Watson couldn't integrate with the hospital's newly adopted Epic EHR

system and struggled to interpret physician notes [20]. At Rigshospitalet in Copenhagen, oncologists found Watson agreed with their clinical judgment in only approximately 33 percent of cases. The hospital declined to purchase the system. In South Korea, Watson's most frequently recommended breast cancer treatments were not covered by the national insurance system [19].

The lessons are precise. Training on synthetic or single-institution data produces recommendations that reflect individual preferences, not generalizable evidence. Global deployment without independent validation exposes patients to unquantifiable risk. Without mandatory adverse event reporting for AI clinical decision support tools, the harms became essentially invisible. IBM sold Watson Health in 2021.

Epic Sepsis Model: When Accuracy Meets Reality

Sepsis kills approximately 350,000 Americans annually, and speed of detection is among the most important determinants of survival. The Epic Sepsis Model (ESM), embedded in the Epic EHR platform used by most American hospitals, promised to help. It calculated a sepsis risk score every 15 minutes for every hospitalized patient, drawing on vital signs, laboratory values, and EHR data. Epic's internal validation reported an area under the receiver operating curve of 0.76 to 0.83.

Then Andrew Wong and colleagues at the University of Michigan conducted an independent external validation. Their study, published in *JAMA Internal Medicine* in 2021, covered 27,697 patients across 38,455 hospitalizations [21]. The ESM's AUC in real-world practice was 0.63, substantially worse than claimed. At the recommended alert threshold, the model's sensitivity was 33 percent. It was missing two out of every three sepsis cases. Its positive predictive value was 12 percent, meaning only about 1 in 8 alerts represented actual sepsis, and it generated alerts for 18 percent of all hospitalizations, creating an alert burden of approximately 6,971 flagged patients in the study cohort [21].

The study identified the training flaw clearly. Epic had used hospital billing codes to define sepsis outcomes, a methodology known to be inaccurate. As Wong stated: "The model is essentially trained to identify which patients will be billed for sepsis, not which patients actually develop the clinical

criteria for sepsis" [21]. The model had learned the patterns of documentation and billing rather than the physiology of infection.

A separate real-world study at another institution found the ESM's positive predictive value at 13 percent, and documented 97,216 sepsis alerts over a two-year period, approximately 133 alerts every day [21]. Dr. Anand Habib at UCSF described the clinical reality: a patient with rapidly progressive interstitial lung disease, short of breath and tachycardic but not septic, was generating ESM alerts two to three times daily, causing nurses significant moral distress as they repeatedly questioned whether they were failing the patient by not treating for sepsis.

Alert fatigue from tools with low positive predictive value—meaning that a high proportion of the alerts it generates are false alarms—is not just an inconvenience. Published studies on sepsis alerts specifically document override rates of 49 to 96 percent [22]. When every alert feels like noise, genuine alarms are ignored at the same rate as irrelevant ones. Life-threatening alerts have been overridden at rates as high as 74 percent in desensitized clinical environments [22]. A 67 percent miss rate for sepsis, combined with an alert system that flags one in five hospitalized patients, is a patient safety failure that the original vendor validation couldn't detect, because that validation never tested the tool in conditions resembling clinical practice.

TruDi Navigation System: AI as Marketing, Patients as Collateral

In 2021, Acclarent, a Johnson and Johnson subsidiary, added a machine learning algorithm to its TruDi Navigation System, a device guiding ear, nose, and throat surgeons through sinuplasty procedures by confirming instrument position in patients' heads. According to a subsequent lawsuit, the AI was integrated primarily as a marketing tool.

Before the AI integration: seven adverse event reports over three years, one involving patient harm. After the AI integration through November 2025: at least 100 malfunctions and adverse incidents, and at least 10 patients injured [23].

The injuries included cerebrospinal fluid leaking from a patient's nose, a surgeon unintentionally puncturing the base of a patient's skull, and two

patients suffering strokes after their carotid arteries were accidentally damaged. Erin Ralph, a patient in Fort Worth, Texas, had a carotid artery injury during a June 2022 sinuplasty. She discovered she had suffered a stroke after leaving the hospital. A portion of her skull was removed to allow her brain room to swell. She required intensive care for five days. In a subsequent interview she stated: “It is hard to walk without a brace and [to] regain full function of my left arm.” The surgeon’s own records reportedly stated he “had no idea [the instrument] was anywhere near the carotid artery” [23].

A second patient, Donna Fernihough, suffered a carotid artery injury in May 2023 during which blood reportedly sprayed across the operating room. Her lawsuit alleged Acclarent had “lowered its safety standards to rush the new technology to market,” targeting only 80 percent accuracy for some components before integration [23].

Dr. Alexander Everhart of Washington University School of Medicine put it plainly: “I believe the FDA’s conventional method of regulating medical devices is insufficient to ensure the safety and effectiveness of AI-enabled technologies. We are relying on manufacturers to ensure quality” [23]. Most AI-enabled medical devices are not required to undergo clinical trials before approval. Manufacturers can reference previously authorized non-AI devices, adding AI features without separate prospective validation studies. The jump from 7 adverse reports before the AI integration to more than 100 afterward is the empirical argument for why this regulatory structure is inadequate.

Risks and Limitations

While we work to improve AI systems, we must understand the risks and limitations of the technologies we deploy. The flaws in real-world AI models create genuine hazards for patients, particularly when clinicians rely on automated recommendations without sufficient scrutiny.

The Automation Bias Trap

When a biased AI is wrong, clinicians tend to agree with it anyway. This is automation bias: the documented tendency to defer to automated systems

even when independent judgment would reach a different conclusion. A multicenter randomized clinical vignette survey study, published in *JAMA* in 2023, tested hospitalist physicians, nurse practitioners, and physician assistants on clinical diagnosis of acute respiratory failure. When shown an accurate AI model, their diagnostic accuracy increased by 4.4 percentage points. When shown a biased AI model, one that systematically misdiagnosed older women with COPD as having pneumonia, their accuracy decreased by 11.3 percentage points [2].

The investigators then asked whether providing AI model explanations could protect against this effect. It couldn't. Even with explainability features active, accuracy still decreased by 9.1 percentage points in the biased AI condition, a statistically nonsignificant change. The conclusion was unambiguous: "AI model explanations did not help clinicians recognize systematically biased AI models" [2]. The transparency features that regulators increasingly encourage as a safeguard were insufficient to protect patients from a biased tool. At the Radiological Society of North America (RSNA) 2024, a multisite prospective study of 220 radiologists and physicians found that when AI advice was incorrect, physician accuracy fell to 23.6 percent, a catastrophic decline driven by overreliance on the machine's erroneous reading [24].

In my experience, this is one of the most under-discussed dangers in clinical AI deployment. We train physicians to trust well-calibrated tools. That trust becomes a liability the moment the tool is wrong.

Alert Fatigue: The Signal Lost in Noise

Clinical decision support systems are overridden at rates of 49 to 96 percent across published studies [22]. Drug-dose alerts are overridden in 82 to 96.8 percent of cases. Drug-drug interaction alerts are overridden at 56 to 95.6 percent [22]. These aren't rates of clinician negligence. They are the rational behavioral adaptation of professionals overwhelmed by a signal-to-noise ratio that renders most alerts meaningless. The Brigham and Women's Hospital experience crystallized the danger: clinicians overrode 73.3 percent of medication alerts, and of those, 40 percent were inappropriately dismissed. Among overridden renal alerts, 12.5 percent of dose-change alerts and 29.6 percent of avoid-medication alerts had genuine clinical

basis, meaning real patients were exposed to real harm because the alert system had destroyed its own credibility [22].

The Dataset Shift Failure Mode

A model validated in one setting fails silently in another. The *New England Journal of Medicine*'s 2021 paper by Finlayson and colleagues documented the taxonomy of this failure: population shift (a model trained on predominantly white patients underperforms on diverse ones), acquisition device shift (an AI pneumonia detector trained on one scanner brand degrades when images come from a different scanner), EHR software shift (an ICU mortality model trained on CareVue data collapsed when the hospital switched to MetaVision), and coding practice shift (the ICD-9 to ICD-10 transition caused more than 20 percent of diagnostic categories to experience instantaneous prevalence shifts) [25].

The Google Health diabetic retinopathy system offers the case study in full detail. The tool achieved greater than 90 percent accuracy against specialists in controlled laboratory settings. Deployed in 11 Thai clinics, more than 20 percent of images were rejected by the AI because nurses captured them in the suboptimal lighting absent from training data. Patients whose scans were rejected had to reschedule at specialist clinics on different days, adding burdens their circumstances often made impossible to meet. Slow Internet connectivity in rural clinics created upload delays. Nurses reported patients waiting since 6 a.m. with fewer than 10 screenings completed in the first two hours [26]. Lab accuracy had not predicted clinical value.

The Validation Gap

A 2025 Health Affairs study of more than 2,400 U.S. hospitals found that only 61 percent of hospitals using AI or predictive models evaluated those models for accuracy using local data [27]. Only 44 percent tested for bias at all [27]. Hospitals serving marginalized patient populations, including rural, critical access, and high social deprivation index facilities, were less likely to use AI and also less likely to evaluate it when they did. The study concluded: "Hospitals relying on EHR developer-made models and less often locally evaluating could be more susceptible to using inaccurate or biased models that could harm patients" [27]. ECRI, the nonprofit patient

safety organization, named AI-enabled health technologies the number-one health technology hazard for 2025, the first time AI had topped the list, citing hallucinations, training data bias, and performance degradation over time as the primary concerns [28].

Regulatory and Legal Landscape

The legal and regulatory environment governing clinical AI has not kept pace with technological innovation. Regulatory agencies, legal frameworks, and professional standards struggle to address the unique challenges posed by machine learning systems in healthcare.

The FDA and the Pace Problem

The FDA has authorized more than 1,016 AI/ML-enabled medical devices as of December 2024. Radiology accounts for 76 percent of cleared devices. Cardiovascular AI represents 9.8 percent [29]. The pathway to market is the 510(k) clearance process, used in 96.4 percent of cases: a pathway designed for devices that demonstrate substantial equivalence to a previously cleared device, not for adaptive algorithms that learn and change after deployment [29].

The FDA issued its final guidance on Predetermined Change Control Plans (PCCPs) in December 2024, a framework allowing manufacturers to specify in advance how their algorithms may change without requiring new regulatory submissions for every modification [29]. The guidance recommends that manufacturers consider the characteristics of intended use populations, including ethnicity and gender, in their change control plans. In January 2025, the FDA published a draft guidance proposing lifecycle management requirements including bias analysis and real-world performance monitoring.

The problem is the gap between aspiration and adoption. As of December 2024, only 15 of 1,016 cleared AI devices (1.5 percent) had reported a PCCP at all. Only 240 devices (23.7 percent) reported any dataset demographic information in their approval summaries [29]. The regulatory infrastructure for adaptive AI exists in draft form. Clinical deployment of AI has not waited for it.

The EU AI Act and Biden to Trump

The European Union's AI Act, which entered into force in August 2024, classifies most commercial AI medical devices as high-risk, requiring explicit bias detection and mitigation requirements, continuous risk management, and human oversight provisions explicitly designed to allow practitioners to recognize and manage automation bias [30]. Violations of prohibited AI uses carry fines up to 35 million euros or 7 percent of global annual turnover. High-risk medical device AI faces full compliance requirements beginning August 2027 [30].

President Biden's Executive Order 14110, signed October 30, 2023, directed HHS to establish an AI Task Force, develop a Strategy on AI Quality, and create an AI Safety Program to capture clinical errors resulting from AI in healthcare settings, with explicit acknowledgment of AI's potential to perpetuate and introduce biases in clinical algorithms [31]. The Trump administration revoked EO 14110 in early 2025, directing agencies to prioritize AI innovation and remove regulatory barriers. In December 2025, a new executive order directed the Justice Department to challenge state AI laws deemed inconsistent with federal policy, specifically citing Colorado's algorithmic bias audit requirement as a target [31].

State-level regulation still represents meaningful patient protection. California AB 3030, signed in September 2024 and effective January 2025, requires healthcare providers using generative AI for patient clinical communications to disclose that the communication was AI-generated and provide instructions for contacting a human provider [32]. Colorado SB 169 prohibits insurers from using algorithms or predictive models that result in unfair discrimination based on race, color, national origin, sex, and other protected characteristics, though this law faces a federal preemption challenge under the December 2025 executive order [32].

Legal Accountability: The Lawsuits Arrive

The legal system has begun catching up to clinical AI deployment in several directions simultaneously. The class action lawsuit *Estate of Lokken v. UnitedHealth Group*, filed in November 2023, alleges that UnitedHealth used its nH Predict algorithm, developed by its naviHealth subsidiary, to systematically deny post-acute care coverage for elderly Medicare

Advantage patients. A STAT News investigation found that UnitedHealth set company targets to keep patient rehabilitation stays within 1 percent of the length projected by the algorithm, with employees facing discipline for missing targets. The algorithm allegedly has a 90 percent error rate on appeals: approximately 90 percent of denied claims that were appealed were overturned by federal administrative law judges. In February 2025, a federal court allowed breach of contract claims to survive a motion to dismiss, finding that plan documents promising physician-led medical review may have been violated by near-automatic AI gatekeeping [33].

Cigna faced similar litigation over its Pxdx algorithm, which allegedly enabled physicians to automatically deny payments in batches without individual review of patient circumstances [34]. Humana faced a parallel class action. All three major Medicare Advantage insurers were simultaneously under legal scrutiny for AI-driven coverage denials, a pattern that prompted the CMS proposed rule of December 2024, which proposed specific AI guardrails for Medicare Advantage plans, though those guardrails were ultimately not included in the April 2025 final rule [34].

A cross-sectional study published in *JAMA Health Forum* in 2025 examined 950 FDA-cleared AI medical devices: 60 (6.3 percent) were associated with 182 recall events, with 43.4 percent of recalls occurring within the first 12 months of clearance, approximately double the recall rate for all 510(k) devices in the same period [35]. Devices lacking clinical validation studies had a mean recall rate of 3.4 per device versus 2.0 for those with prospective validation studies [35].

Ethical Implications

Beyond the technical and regulatory challenges, clinical AI deployment raises profound ethical questions. These questions go beyond whether an algorithm works well on average. They require us to ask whether the systems we deploy serve human dignity and health equity.

Algorithmic Justice: Beyond Fairness

There is a real distinction between algorithmic fairness and algorithmic justice, and healthcare leaders need to understand both. Fairness asks

whether an AI tool performs equally well across demographic groups. Justice asks the harder question: whether a tool is designed in a way that actively confronts, rather than merely stops propagating, historical harm.

The UCSF Algorithmic Justice Workshop, which convened leading researchers and clinicians and issued its report in July 2024, defined three core domains: transparency and explainability of algorithms, community engagement throughout the development and deployment lifecycle, and accountability that operates against explicit equity and justice standards [36]. The workshop identified what scholars have called “bounded justice,” the recognition that technical fairness metrics treat symptoms rather than structural causes. Optimizing a model for demographic parity may still leave intact all the structural barriers that generated health disparities in the first place.

The American Medical Association’s Ethics-Evidence-Equity framework articulates the same imperative in clinical governance terms [37]. Trustworthy AI must simultaneously uphold scientific standards (Evidence), align with the values of medicine (Ethics), and apply an equity lens at every stage of design, validation, deployment, and monitoring (Equity). The AMA has been explicit: “An equity lens must be applied intentionally from the very beginning, in populating the design and testing team, the framing of the problem to be solved, the training data set selected, and the design and evaluation of the algorithm itself” [37].

The WHO’s six ethical principles for AI in health, developed through an 18-month international consultation and published in 2021, provide the globally applicable framework: protecting human autonomy, promoting well-being and safety, ensuring transparency and explainability, building responsibility and accountability, ensuring inclusiveness and equity, and promoting AI that is responsive and sustainable [38]. The WHO explicitly warned against “data colonialism,” in which high-income countries collect training data from low- and middle-income settings without parallel protections, and stated that existing biases in healthcare services based on race, ethnicity, age, and gender that are encoded in training algorithms “must be overcome” [38].

Informed Consent in the Age of Algorithmic Medicine

When a physician uses an AI tool to assist a diagnosis, does the patient have a right to know? *JAMA* published guidance in July 2025 asserting that the legal doctrine of informed consent "requires disclosing information that is material to a reasonable patient's decision to accept health care" and that AI use may meet this threshold for high-impact decisions [39]. A risk-stratified approach is taking shape in the literature: notification for confirmatory diagnostics; education about algorithm basis and limitations for prognostic tools influencing care plans; and full informed consent, including disclosure of known biases and demographic performance disparities, for life-altering decisions such as transplant candidacy or end-of-life planning [39].

The fairness metric impossibility theorem is a mathematical reality that consent processes need to address. Demographic parity and calibration can't be simultaneously satisfied when base rates differ across demographic groups, a formal proof with direct clinical implications [40]. Every fairness criterion represents a trade-off, and the choice of which criterion a tool is optimized for is a values decision, not just a technical one. Patients, particularly those from groups that have historically been disadvantaged, have a legitimate interest in knowing which trade-offs were made on their behalf.

What Successful Debiasing Looks Like

The record is not only of failure. After Obermeyer and colleagues published their *Science* paper, the algorithm manufacturer reviewed the findings and initiated product revision. Recalibrating the objective function, replacing healthcare costs with direct health indicators such as chronic condition counts, eliminated the racial bias and increased Black patient enrollment from 17.7 percent to 46.5 percent [1]. The Daneshjou study showed that when AI dermatology models were fine-tuned on the diverse Ddi dataset, the performance gap between light and dark skin tones was eliminated. Fine-tuned models actually outperformed dermatologists on darker skin tones [3]. The eGFR race correction, which had delayed nephrology referrals and transplant eligibility for Black patients for decades, was removed from clinical calculators following advocacy by the American Society of Nephrology and the National Kidney Foundation [41]. New York City's Coalition to End Racism in Clinical Algorithms (CERCA)

successfully engaged seven of nine major NYC health systems to remove multiple race-based algorithmic modifiers from clinical decision support tools [42].

These interventions identified the proxy variable carrying the bias: cost, skin tone training data, race as a physiological variable. They replaced it with a more direct measure. They required institutional commitment, multidisciplinary teams, and clinical education alongside the technical fix. The eGFR case is instructive precisely because the technical fix alone was insufficient: removing the race variable from the equation did not change nephrology referral rates, because structural barriers to care persisted beyond the algorithm [41]. Technical de-biasing must be accompanied by structural intervention.

Leadership Implementation Framework

The same year that ECRI named AI the top health technology hazard, a survey found that only 16 percent of healthcare organizations had system-wide AI governance frameworks in place, while the AI healthcare market expanded by 42 percent [43]. Nearly half of healthcare organizations lacked any approval processes specific to AI technologies. The gap between deployment velocity and governance maturity is where patients are harmed.

AMA Governance for Automated Intelligence

The AMA's Governance for Augmented Intelligence toolkit provides a ready operational framework: establish executive accountability and a governance structure with CEO and board-level commitment; form a working group with defined priorities, processes, and policies; assess current AI activity and establish priorities; develop AI policies with minimum required elements; define project intake, vendor evaluation, and assessment processes; update standard planning and implementation processes; establish oversight and monitoring; and support organizational AI readiness through stakeholder engagement [37]. Each step is a structured institutional action, not an aspiration.

Coalition for Health AI

The Coalition for Health AI (CHAI), a provider-led coalition of more than 3,000 organizations including health systems, patient advocacy groups, academia, government agencies, and industry, published its Blueprint for Trustworthy AI in April 2023, followed by a 187-page Assurance Standards Guide with reporting checklists in June 2024 [44]. The guide provides a six-step AI lifecycle framework with assurance reporting checklists applicable to the most common clinical AI deployment categories, including predictive EHR risk, imaging diagnostics, and prior authorization. The CHAI framework aligns with the National Academy of Medicine's AI Code of Conduct, the White House Blueprint for an AI Bill of Rights, and NIST frameworks [44].

Standing Together

The Standing Together consensus recommendations, published in *NEJM AI* in 2024, represent the most thorough international statement on AI bias standards in healthcare to date. More than 350 representatives from 58 countries contributed through a Delphi process and public consultation, establishing requirements for transparent reporting of dataset characteristics, demographic representation, and algorithmic performance stratified by demographic subgroups [45]. The consensus establishes what mandatory bias reporting for AI medical devices should look like before deployment.

Model Cards and Datasheets for Datasets

Model Cards and Datasheets for Datasets, tools originally developed at Google Research and codified for healthcare contexts, provide the documentation infrastructure [46]. Model Cards are short, standardized documents accompanying trained models that include benchmarked evaluation across demographic subgroups, intersectional performance measures (age and race, not just each separately), and known limitations. Datasheets for datasets document the composition, collection process, and known biases of every training dataset. Together they provide the transparency infrastructure without which procurement teams can't make

informed decisions about algorithmic safety. Healthcare organizations should require these documents from every AI vendor before contract.

The Importance of Diverse Development Teams

Building diverse AI development teams is not just a social good. It is a bias mitigation strategy with documented effectiveness. Less than 25 percent of AI development staff currently identify as racial or ethnic minorities. Fewer than one-third of companies report that their AI teams are racially diverse [47]. Diverse teams are more likely to flag underrepresentation in training data, to prioritize minimizing performance disparities over maximizing aggregate accuracy, and to design for the clinical realities of under-resourced settings rather than academic medical centers. The EU AI Act explicitly requires diverse development teams. U.S. health systems should adopt this standard proactively.

Community engagement remains rare in healthcare AI tool development [48]. The UCSF Algorithmic Justice framework requires "reliable and consequential community engagement at all stages and facets of the development and deployment of institutional algorithms," a standard that requires patient and community advisory boards with genuine authority, not ceremonial participation [36].

Future Outlook

The institutional machinery for safer healthcare AI is being assembled, if not yet quickly enough.

> **The NIST AI Risk Management Framework**, published in January 2023 and incorporated by reference in Biden EO 14110, provides a voluntary but increasingly authoritative four-function structure: Govern, Map, Measure, Manage, with explicit healthcare application and explicit inclusion of bias in training data as a risk to be measured and managed [49]. The Trump AI Action Plan of July 2025 directed NIST to launch sector-specific initiatives in healthcare to develop national AI standards, a bipartisan recognition that the framework's voluntary character must eventually acquire regulatory force.

The NIH All of Us Research Program has enrolled more than 832,000 participants, with 80 percent drawn from groups underrepresented in biomedical research by race, ethnicity, age, geography, sexual and gender identity, income, education, and disability [50]. The program holds more than 452,000 electronic health records, over 245,000 whole-genome sequences, and approximately 1.96 million surveys: resources that did not exist five years ago. More than 12,475 registered researchers from over 865 organizations, including 44 historically Black colleges and universities and 73 Hispanic-Serving Institutions, have accessed the data [50]. This is what representative training data can look like at scale.

The FDA's Draft Guidance on AI-Enabled Device Software Functions, published in January 2025, proposes lifecycle management requirements including bias analysis and real-world performance monitoring, the closest the agency has come to binding post-deployment obligations for AI medical devices [29]. The guidance remains in draft form and its implementation trajectory under the current administration is uncertain. Its publication represents regulatory acknowledgment that pre-market approval is a necessary but not sufficient condition for patient safety.

Continuous monitoring and drift detection (the process of identifying when model performance deteriorates over time) are moving from research concepts to clinical infrastructure. Nestor and colleagues demonstrated in a study of 143,049 patients across seven Toronto hospitals that proactive monitoring with drift-triggered model updating improved performance during the COVID-19 pandemic, showing that real-world model degradation is detectable and correctable in prospective clinical settings [51]. Daneshjou and colleagues showed that fine-tuning on diverse images eliminated the skin tone performance gap, confirming that the technical solutions are available [3]. The constraint is not knowledge. It is implementation.

The future of equitable healthcare AI depends on a set of institutional commitments that have nothing to do with the sophistication of the models: transparent training data documentation, stratified performance reporting, continuous post-deployment monitoring, diverse development teams,

community engagement, and governance structures with genuine accountability. The evidence on what goes wrong, and what makes it right, is already in hand. What changes outcomes is whether leaders choose to act on it.

Key Takeaways

Algorithmic bias is not a future risk; it is an active patient safety crisis. The Obermeyer study documented Black patient enrollment in care management programs at less than half the appropriate rate. The pulse oximetry data showed Black patients waited 278.5 additional minutes for supplemental oxygen during COVID-19. These are population-scale harms occurring now.

Biased proxy variables are the most common mechanism. When an algorithm uses healthcare costs to predict health needs or builds dermatology AI on light-skinned training data, it doesn't need to include race as a variable to produce racially disparate outcomes. Leaders must ask: what is this model's outcome variable, and does it inherit structural inequity?

Internal validation by developers substantially overstates real-world performance. The Epic Sepsis Model's AUC fell from 0.76–0.83 (internal) to 0.63 (external validation). IBM Watson was deployed at 230 hospitals before any independent validation existed. Independent external validation should be a precondition for deployment, not an optional supplement.

Automation bias converts AI errors into compounded patient harm. A biased AI reduced clinician diagnostic accuracy by 11.3 percentage points even when explanations were provided. Clinicians must be trained to challenge AI outputs, not simply accept them, particularly for demographic groups with known performance disparities.

Alert fatigue is a governance failure, not a clinician failure. Override rates of 49 to 96 percent are the rational behavioral response to tools with low positive predictive value deployed at scale. The

signal-to-noise ratio of AI alerts determines whether the tool enhances or degrades patient safety.

Dataset shift silently degrades deployed models. Every EHR upgrade, coding transition, population demographic shift, or hardware change creates conditions for AI performance degradation without any visible warning. Continuous post-deployment monitoring is not optional, it is the mechanism that makes deployment safe.

Technical debiasing requires structural accompaniment. Removing the race variable from the eGFR equation did not change nephrology referral rates. Algorithmic corrections must be paired with clinical education, change management, and structural interventions in care delivery.

Sixty-one percent of hospitals test AI for accuracy; only 44 percent test for bias. Hospitals serving the most vulnerable populations are least likely to do either. Health systems must treat bias audits as standing governance obligations, not one-time activities.

Regulatory frameworks exist; adoption is the gap. The AMA eight-step governance toolkit, CHAI Assurance Standards, Standing Together consensus, NIST AI RMF, and WHO's six guiding principles collectively provide actionable governance infrastructure. Only 16 percent of health systems had system-wide AI governance frameworks in 2024. The frameworks are not the constraint.

Equitable AI is achievable. Recalibrating the Optum algorithm eliminated the racial enrollment gap and tripled Black patient access to care management. Fine-tuning dermatology AI on diverse images eliminated skin tone performance disparities. The tools work better for everyone when they're built for everyone. Equity and performance are not in tension, they're the same goal.

References

1. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019; 366(6464):447-453. doi:10.1126/science.aax2342.

2. Jabbour S, Fouhey D, Shepard S, Valley TS, Kazerooni EA, Banovic N, et al. Measuring the Impact of AI in the Diagnosis of Hospitalized Patients: A Randomized Clinical Vignette Survey Study. JAMA. 2023; 330(23):2275-2284. doi:10.1001/jama.2023.22295.

3. Daneshjou R, Vodrahalli K, Novoa RA, Jenkins M, Liang W, Rotemberg V, et al. Disparities in dermatology AI performance on a diverse, curated clinical image set. Sci Adv. 2022; 8(31):eabq6147. doi:10.1126/sciadv.abq6147. PMID 35960806.

4. ProBiologists. Dermatology AI geographic diversity [Internet]. 2022 [cited 2026 Jul 4]. Available from: https://www.probiologists.com.

5. Groh M, Harris C, Soenksen L, Lau F, Han R, Kim A, et al. Evaluating deep neural networks trained on clinical images in dermatology with the Fitzpatrick 17k dataset. Proc IEEE/CVF Conf Comput Vis Pattern Recognit Workshops (CVPRW). 2021;1820-1828. doi:10.1109/CVPRW53098.2021.00201.

6. Tobin MJ, Jubran A. Pulse oximetry, racial bias and statistical bias. Ann Intensive Care. 2022; 12(1):2. doi:10.1186/s13613-021-00974-7.

7. Fawzy A, Wu TD, Wang K, Robinson ML, Farha J, Bradke A, et al. Racial and ethnic discrepancy in pulse oximetry and delayed identification of treatment eligibility among patients with COVID-19. JAMA Intern Med. 2022; 182(7):730-738. doi:10.1001/jamainternmed.2022.1906. PMID 35639368.

8. Sudat SEK, Wesson P, Rhoads KF, Brown S, Aboelata N, Pressman AR, et al. Racial disparities in pulse oximeter device inaccuracy and estimated clinical impact on COVID-19 treatment course. Am J Epidemiol. 2023; 192(5):703-713. doi:10.1093/aje/kwac164.

9. Wong AKI, Charpignon M, Kim H, Josef C, de Hond AAH, Fojas JJ, et al. Analysis of discrepancies between pulse oximetry and arterial oxygen saturation measurements by race and ethnicity and association with organ dysfunction and mortality. JAMA Netw Open. 2021; 4(11):e2131674. doi:10.1001/jamanetworkopen.2021.31674.

10. Shi C, Goodall M, Dumville J, Hill J, Norman G, Hamer O, et al. The accuracy of pulse oximetry in measuring oxygen saturation by levels of skin pigmentation: a systematic review and meta-analysis. BMC Med. 2022; 20(1):267. doi:10.1186/s12916-022-02452-8.

11. Straw I, Rees G, Nachev P. Sex-based performance disparities in machine learning algorithms for cardiac disease prediction: exploratory study. J Med Internet Res. 2024; 26:e46936. doi:10.2196/46936.

12. Krieger K, Hameed I, Quer G, Mack C, Savic M, Mantaj P, et al. GPT4 reinforces historical gender bias in diagnosing cardiovascular symptoms in women. Eur Heart J. 2025; 46(Suppl 1):ehaf784.4497. doi:10.1093/eurheartj/ehaf784.4497.

13. de Andrade JBC, de Medeiros Cavalcante MAO, Lopes TLM, Treigher JMS, Balsells MD, Vasconcelos JL, et al. Discovery of data quality issues in electronic health records: profound consequences for critical care medicine applications – a systematized review. Crit Care. 2026;30:XX. doi:10.1186/s13054-025-05677-0. PMID 41508097.

14. World Health Organization. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models. Geneva: WHO; 2024. Available from: https://www.who.int/publications/i/item/9789240084759.

15. Otokiti AU, Shih HJ, Williams KS. Gender and racial bias unveiled: clinical artificial intelligence (AI) and machine learning (ML) algorithms are fanning the flames of inequity. Oxford Open Digit Health. 2025; 1(1):oqaf027. doi:10.1093/oodh/oqaf027. PMID 41163810.

16. Muralidharan V, Adewale BA, Huang CJ, Nta MT, Ademiju PO, Pathmarajah P, et al. A scoping review of reporting gaps in FDA-approved AI medical devices. npj Digit Med. 2024; 7(1):273. doi:10.1038/s41746-024-01270-x.

17. Gichoya JW, Banerjee I, Bhimireddy AR, Burns JL, Celi LA, Chen LC, et al. AI recognition of patient race in medical imaging: a modelling study. Lancet Digit Health. 2022; 4(6):e406-e414. doi:10.1016/S2589-7500(22)00063-2.

18. Velichkovska B, Gjoreski H, Denkovski D, Kalendar M, Mullan ID, Wawira Gichoya J, et al. Evaluation frameworks for clinical AI systems. BMJ Health Care Inform. 2025; 32(1):e101098. doi:10.1136/bmjhci-2024-101098. PMID 40639839.

19. Ross C, Swetlitz I. IBM Watson's questionable clinical impact. STAT. 2018. Available from: https://www.statnews.com/2018/07/25/ibm-watson-recommended-unsafe-incorrect-treatments. Accessed 2026.

20. Herper M. MD Anderson halts IBM Watson project. Forbes. 2017. Available from: https://www.forbes.com/sites/matthewherper/2017/02/19/md-anderson-benches-ibm-watson-in-setback-for-artificial-intelligence-in-medicine. Accessed 2026.

21. Wong A, Otles E, Donnelly JP, Krumm A, McCullough J, DeTroyer-Cooley O, et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Intern Med. 2021; 181(8):1065-1070. doi:10.1001/jamainternmed.2021.2626.

22. Ginestra JC, Giannini HM, Schweickert WD, Meadows L, Lynch MJ, Pavan K, et al. Clinician perception of a machine learning-based early warning system designed to predict severe sepsis and septic shock. Crit Care Med. 2019; 47(11):1477-1484. doi:10.1097/CCM.0000000000003803.

23. Reuters. AI enters clinical decision-making: investigation [Internet]. 2026 [cited 2026 Jul 4]. Available from: https://www.reuters.com/investigations.

24. Prinster D, Mahmood A, Saria S, Jeudy J, Lin CT, Yi PH, et al. Care to explain? AI explanation types differentially impact chest radiograph diagnostic performance and physician trust in AI. Radiology. 2024; 313(2):e233261. doi:10.1148/radiol.233261. PMID 39560483.

25. Finlayson SG, Subbaswamy A, Singh K, Bowers J, Kupke A, Zittrain J, et al. The clinician and dataset shift in artificial intelligence. N Engl J Med. 2021; 385(3):283-286. doi:10.1056/NEJMc2104626.

26. Beede E, Baylor E, Hersch F, Iurchenko A, Wilcox L, Ruamviboonsuk P, et al. A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York: ACM; 2020. p. 1-12. doi:10.1145/3313831.3376718.

27. Nong P, Adler-Milstein J, Apathy NC, Holmgren AJ, Everson J. Current use and evaluation of artificial intelligence and predictive models in US hospitals. Health Aff (Millwood). 2025; 44(1):90-98. doi:10.1377/hlthaff.2024.00842.

28. ECRI. Top 10 health technology hazards for 2024. Plymouth Meeting (PA): ECRI; 2024.

29. US Food and Drug Administration. Artificial intelligence/machine learning (AI/ML)-based software as a medical device (SaMD) action plan. Silver Spring (MD): FDA; 2021 Jan. Available from: https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device.

30. European Parliament. Artificial Intelligence Act (Regulation (EU) 2024/1689). 2024. Available from: https://artificialintelligenceact.eu. Accessed 2026.

31. The White House. Executive Order 14110 on safe, secure, and trustworthy AI. 2023.

32. California Legislature. Artificial intelligence legislation: AB 3030 (Health care services: artificial intelligence, Chapter 848, Statutes of 2024); SB 1120 (Physicians Make Decisions Act, Chapter 879, Statutes of 2024). Available from: https://legiscan.com/CA/text/SB1120/id/3023335.

33. Lokken v UnitedHealth Group. Court filing. 2023.

34. Cigna PxDx litigation. 2023.

35. Lee B, Kramer P, Sandri S, Chanda R, Favorito C, Nasef O, et al. Early recalls and clinical validation gaps in artificial intelligence-enabled

medical devices. JAMA Health Forum. 2025; 6(8):e253172. doi:10.1001/jamahealthforum.2025.3172.

36. University of California San Francisco. Algorithmic justice workshop report. 2024.

37. American Medical Association. Augmented intelligence in health care toolkit. 2025.

38. World Health Organization. Ethics and governance of artificial intelligence for health. Geneva: WHO; 2021.

39. Mello MM, Char D, Xu SH. Physicians' obligations to inform patients about clinical artificial intelligence. JAMA. 2025; 334(9):777-778. doi:10.1001/jama.2025.11417. PMID 40690211.

40. Chouldechova A. Fair prediction with disparate impact: a study of bias in recidivism prediction instruments. Big Data. 2017; 5(2):153-163. doi:10.1089/big.2016.0047.

41. Stanford Institute for Human-Centered AI. The complexities of race adjustment in health algorithms [policy brief]. Stanford (CA): Stanford HAI; 2024. Available from: https://hai.stanford.edu/assets/files/policy-brief-race-adjustment-in-health-algorithms.pdf.

42. Digital Medicine Society. The playbook: implementing AI in healthcare. Boston (MA): DiMe; 2025. Available from: https://dimesociety.org/ai-implementation-in-healthcare-playbook.

43. Censinet. Healthcare AI governance survey report. 2025.

44. Coalition for Health AI. Blueprint for trustworthy AI. 2024.

45. Alderman JE, Palmer J, Laws E, McCradden MD, Ordish J, Ghassemi M, et al. Tackling algorithmic bias and promoting transparency in health datasets: the STANDING together consensus recommendations. NEJM AI. 2025; 2(1):AIp2401088. doi:10.1056/AIp2401088.

46. Mitchell M, Wu S, Zaldivar A, Barnes P, Vasserman L, Hutchinson B, et al. Model cards for model reporting. Proc Conf Fairness Accountability Transparency. 2019;220–9.

47. Abbasgholizadeh Rahimi S, Shrivastava R, Brown-Johnson A, Caidor P, Davies C, Idrissi Janati A, et al. EDAI framework for integrating equity, diversity, and inclusion throughout the lifecycle of AI to improve health and oral health care: qualitative study. J Med Internet Res. 2024; 26:e63356. doi:10.2196/63356. PMID 39546793.

48. Loftus TJ, Balch JA, Abbott KL, Hu D, Ruppert MM, Shickel B, Ozrazgat-Baslanti T, Efron PA, Tighe PJ, Hogan WR, Rashidi P, Cardel MI, Upchurch GR Jr, Bihorac A. Community-engaged artificial intelligence research: a scoping review. PLOS Digit Health. 2024; 3(8):e0000561. doi:10.1371/journal.pdig.0000561. PMID: 39178307.

49. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). Gaithersburg (MD): NIST; 2023.

50. Bianchi DW, Brennan PF, Chiang MF, Criswell LA, D'Souza RN, Gibbons GH, et al. The All of Us research program is an opportunity to enhance the diversity of US biomedical research. Nat Med. 2024; 30(2):330-333. doi:10.1038/s41591-023-02744-3. PMID 38374344.

51. Duckworth C, Chmiel FP, Burns DK, Zlatev ZD, White NM, Daniels TWV. Using explainable machine learning to characterise data drift and detect emergent health risks for emergency department admissions during COVID-19. Sci Rep. 2021; 11(1):23017. doi:10.1038/s41598-021-02481-y.

AI in Emergency Medicine: What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution, First Edition. Harvey Castro.
© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

4
Predictive Diagnosis and the End of Missed Symptoms

The one that haunts me is the woman with the headache.

She came into my ER on a Tuesday night. Mid-forties, otherwise healthy, complaining of the worst headache of her life. She'd been to urgent care that morning. They gave her Ibuprofen and sent her home. By evening the pain was unbearable, so she came to us.

I've seen thousands of headaches. The vast majority are benign. Tension, migraine, sinusitis, dehydration. You learn the patterns, and you learn the red flags: sudden onset, worst-ever severity, neck stiffness, neurological changes. This patient had two of four. In a busy ER with a waiting room full of patients, two of four is enough to order a CT. It's also easy enough to rationalize skipping it, especially at 11 p.m. when the scanner queue is long and the radiologist is covering three hospitals.

I ordered the CT. The subarachnoid hemorrhage lit up the screen like a flare. She was in the OR within the hour. She lived. But she lived because I happened to catch the pattern at that moment, on that shift, at that level of fatigue. A different night, a different doctor, a slightly longer queue, and the decision tree bends a different way.

That's the reality the statistics in this chapter describe. The 795,000 Americans who suffer death or permanent disability from diagnostic errors every year aren't abstractions. They're the woman with the headache who went to urgent care first and was told to take Ibuprofen. They're the spinal abscess misdiagnosed 62 percent of the time. They're the subtle fracture on the X-ray that a tired radiologist reads as normal at 2 a.m.

I don't tell this story to indict clinicians. I am one. I've missed things too. The human brain under time pressure, fatigue, and cognitive load has limits that no amount of training fully overcomes. What AI offers in diagnostic medicine is not replacement of the physician's pattern recognition. It's a second set of eyes that doesn't get tired, doesn't get hungry, and doesn't

have a waiting room shaping its attention. This chapter examines what those eyes can see, what they miss, and what the evidence says about how physician and machine together catch what neither catches alone.

Executive Summary

Every year in the United States, approximately 795,000 people die or are left permanently disabled because a diagnosis was missed, delayed, or wrong. That figure, drawn from a 2024 study in *BMJ Quality and Safety*, is not a theoretical estimate of system fragility. It is a measured accounting of preventable catastrophe [1]. The conditions most often missed read like a roster of medicine's most lethal threats: stroke, sepsis, venous thromboembolism, lung cancer, aortic dissection. These are not obscure diagnoses. They are conditions that emergency physicians, primary-care doctors, and specialists encounter daily, and yet the weighted mean diagnostic error rate across what researchers call the "Big Three," vascular events, infections, and cancers, sits at 11.1 percent [1]. In outpatient settings alone, diagnostic errors affect roughly 12 million Americans every year [2].

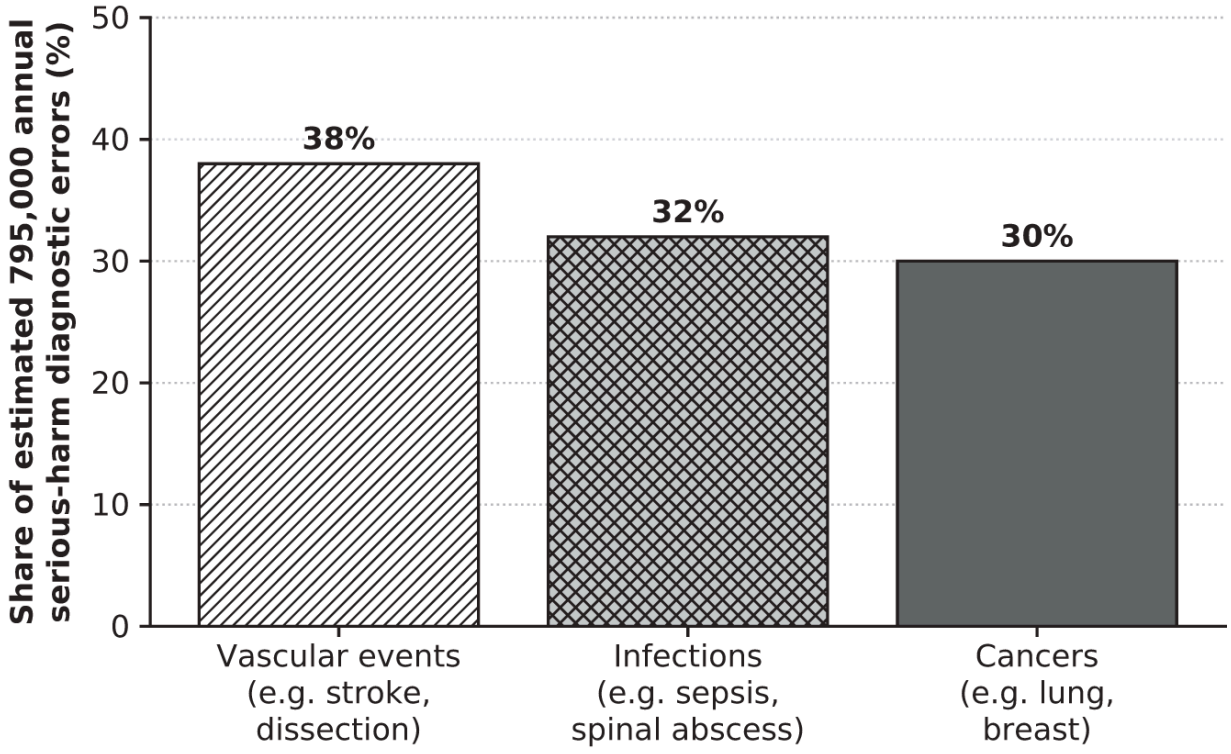
The consequences accumulate quietly: the patient with a 6 mm breast tumor whose mammogram was read as normal, the man with chest pain sent home whose chest X-ray hid a lung cancer in the perihilar shadow, the woman in the intensive-care unit (ICU) whose early sepsis was missed until her organs began to fail. These are not failures of caring. They are, in large part, failures of perception, cognitive bandwidth, and the biological limits of human pattern recognition under conditions of time pressure, incomplete information, and overwhelming volume.

Artificial intelligence doesn't get tired. It doesn't experience anchoring bias after seeing 10 normal scans in a row. It doesn't miss the subtle electrocardiogram (ECG) signature of atrial fibrillation during sinus rhythm or forget to check whether a chest radiograph nodule has grown since the last scan. At its best, AI complements the physician: detecting signals that humans can't reliably see, flagging risk before symptoms appear, and compressing the distance between biological onset and clinical recognition.

This chapter documents the evidence for that promise. It examines where AI predictive diagnostics have been validated in randomized controlled

trials, where they have demonstrated real-world mortality benefit, and where they have already saved specific lives that would otherwise have been lost. It also examines the risks: overdiagnosis, false positives, racial disparities, generalizability failures, and the legal and ethical terrain that healthcare leaders must navigate. A practical framework for the executives, clinicians, and administrators who will decide whether, and how, to bring these tools into their institutions follows.

The stakes are not abstract. Every day that passes without deploying validated AI diagnostic tools in settings where they have proven benefit is a day when preventable deaths continue at a rate that no other industry in the world would accept as normal (see Figure 4.1).



Source: Newman-Toker et al., BMJ Quality & Safety (2024).
Category shares illustrative; total ≈795,000 annual serious harms (range 598k–1,023k).

Figure 4.1 The AI-assisted radiology workflow: scan, AI analysis, risk score, radiologist review

Current Evidence

Diagnostic error remains one of the most persistent and consequential failures in modern medicine. Despite decades of quality improvement work, the data on missed, delayed, and incorrect diagnoses continues to paint a sobering picture. The evidence presented in this section draws from large-scale studies, real-time clinical audits, and AI-assisted diagnostic trials to quantify both the scope of the problem and the measurable impact of emerging solutions.

The Diagnostic Error Crisis

Diagnostic errors remain one of medicine's most stubborn problems, responsible for an estimated 795,000 Americans permanently disabled or killed each year. AI tools targeting this crisis have produced some of the field's most striking results and some of its most cautionary data.

Before examining AI solutions, the scale of the problem demands honest confrontation. The 795,000 annual deaths and permanent disabilities from diagnostic error in the United States come from a rigorous cross-sectional analysis of 21.5 million sampled hospital discharges combined with national cancer registry data, published in *BMJ Quality and Safety* in 2024 by Newman-Toker and colleagues at Johns Hopkins [1]. That number carries a plausible range of 598,000 to 1,023,000 serious harms per year. In a single disease category, venous thromboembolism, the diagnostic error rate reaches 20.4 percent. For aortic aneurysm and dissection, it climbs to 35.6 percent. For spinal abscess, a condition where delayed diagnosis often means permanent paralysis, the error rate is 62.1 percent [1].

Outpatient medicine offers no refuge. A 2014 study in *BMJ Quality and Safety* estimated that diagnostic errors affect approximately 5.08 percent of U.S. adult outpatients per year, translating to roughly 12 million Americans annually, with an estimated half of those errors carrying potential for harm [2]. In the emergency department, the environment where I spent my career as a physician, often making critical diagnoses under time pressure with incomplete information, fractures alone constitute 24 percent of all diagnostic errors. The most missed are not exotic injuries but hand, wrist, and foot fractures that every emergency physician sees daily [3].

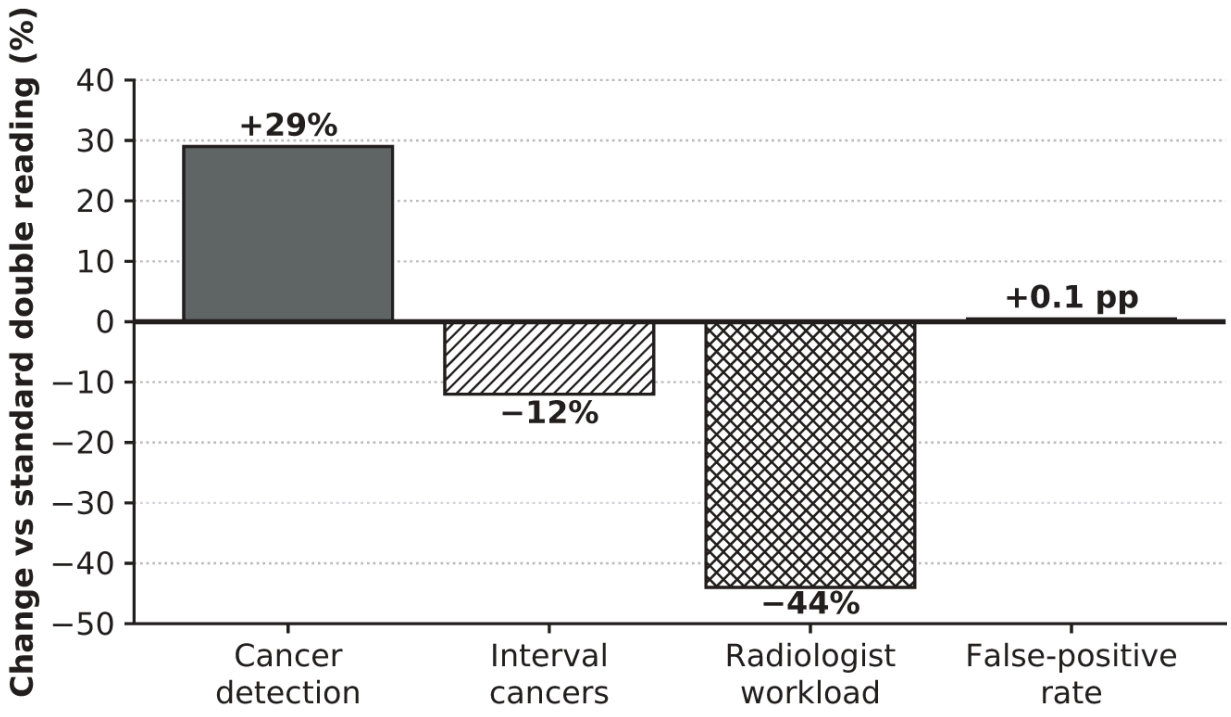
A 2024 single-center study in *BMJ Quality and Safety* found that harmful diagnostic errors occurred in 1 in every 14 general medicine hospital patients, and that 85 percent of those errors were preventable [4]. The most frequently missed diagnoses: heart failure, acute kidney failure, sepsis, and pneumonia, the same quartet that fills hospital wards every night across America.

AI in Cancer Detection

The most compelling early evidence for AI in predictive diagnostics comes from oncology screening, where the biological window between detectable abnormality and clinical symptoms creates a natural opportunity for intervention.

Breast cancer screening has become the proving ground for AI-assisted radiology (see Figure 4.2). The MASAI trial, the first and largest randomized controlled trial of AI in any national cancer screening program, enrolled 105,915 Swedish women and compared AI-supported mammography screening against standard double reading by two radiologists [5]. The final results, published in *The Lancet* in January 2026, showed that AI-supported screening detected 29 percent more cancers than standard double reading, reduced interval cancers by 12 percent, and achieved this while cutting radiologist screen-reading workload by 44 percent [5]. The false positive rate was not increased: 1.5 percent in the AI arm versus 1.4 percent in standard care. The AI group had 16 percent fewer invasive cancers, 21 percent fewer large cancers, and 27 percent fewer aggressive-subtype cancers in the two years following screening, outcomes that point toward genuine clinical benefit, not just earlier labeling.

The German PRAIM trial, published in *Nature Medicine* in 2025, extended these findings to 463,094 women across 12 sites and 119 radiologists. AI-supported screening increased the breast cancer detection rate from 5.7 to 6.7 per 1,000 women, a relative increase of 17.6 percent, without significantly raising recall rates [6].



Source: MASAI randomised controlled trial, Sweden. Manuscript text Chapter 4.

Figure 4.2 Standard double-reading versus AI-assisted single-reading mammography pathways

For lung cancer, the MIT-developed Sybil model takes a different approach: predicting future lung cancer risk from a single low-dose CT scan, without requiring clinical data or radiologist annotations. Trained on 28,162 LDCTs from the National Lung Screening Trial and validated on multiple independent datasets, Sybil achieved a one-year prediction AUC of 0.92 on a held-out NLST cohort, 0.86 at Massachusetts General Hospital, and 0.94 at Chang Gung Memorial Hospital in Taiwan [7]. In a 2025 diversity validation study at the University of Illinois Chicago, a population that was 62 percent non-Hispanic Black, the model maintained strong performance across racial groups, with a one-year AUC of 0.94 [8]. This matters enormously: a diagnostic AI that falls short in diverse populations amplifies health inequity rather than reducing it.

Multicancer early detection represents perhaps the most ambitious application of AI in screening. The Galleri blood test from GRAIL uses cell-free DNA methylation profiling analyzed by machine learning to detect signals from more than 50 cancer types in a single blood draw. The Pathfinder 2 study enrolled 35,878 adults aged 50 or older across 32 U.S.

healthcare systems and reported that adding Galleri to current United States Preventive Services Task Force (USPSTF)-recommended screenings produced a more than sevenfold increase in cancers detected within one year, with 53.5 percent of Galleri-detected cancers at Stage I or II [9]. Approximately 75 percent of the cancers detected had no current recommended screening test. The NHS-Galleri trial, a 142,000-participant randomized controlled trial and the largest multicancer early detection (MCED) trial ever conducted, reported in February 2026 a fourfold higher cancer detection rate in the screened arm and a greater than 20 percent reduction in Stage IV cancer diagnoses in the second and third screening rounds for 12 prespecified deadly cancers [10]. The primary endpoint of statistically significant Stage III-IV reduction was not met, but the trial's secondary findings represent meaningful early evidence that population-level AI-assisted cancer detection is biologically real.

AI in Cardiac Prediction

Cardiovascular disease generates some of the most mature AI diagnostic evidence. The Mayo Clinic's 2019 Lancet study trained a convolutional neural network on 649,931 normal sinus rhythm ECGs from 180,922 patients to detect the "signature" of atrial fibrillation even when the heart was beating normally at the time of the test [11]. The system achieved an AUC of 0.87, with sensitivity of 79.0 percent and specificity of 79.5 percent. When all ECGs from the first month of observation were combined, performance rose to an AUC of 0.90. This was the first demonstration that AI could find preclinical atrial fibrillation from a standard 10-second ECG, opening the possibility of detecting a condition that causes one in four strokes before the first stroke occurs.

Apple Watch AFib detection evidence has now been validated at remarkable scale. A meta-analysis published in *JACC: Advances* in January 2025, pooling 11 studies comprising 4,241 participants, found the Apple Watch ECG algorithm achieved pooled sensitivity of 94.8 percent and specificity of 95 percent for atrial fibrillation, with an AUC of 0.96 [12]. The 2026 EQUAL randomized trial at Amsterdam University Medical Centre assigned 437 patients aged 65 or older with elevated stroke risk to Apple Watch monitoring or standard care, and found a 4.40-fold higher rate of new AFib detection in the smartwatch arm, with 57.1 percent of

smartwatch-detected AFib episodes asymptomatic. Every case in the control arm was detected only after symptoms appeared [13].

AI in Sepsis Detection

Sepsis kills approximately 350,000 Americans every year, and every hour of delayed treatment increases mortality by 4 to 8 percent. The Johns Hopkins TREWS system, the Targeted Real-Time Early Warning System developed by Suchi Saria and colleagues, was prospectively studied across five hospitals in 590,736 patients over two years [14]. For sepsis patients whose TREWS alert was confirmed by a provider within three hours, the system delivered an 18.7 percent relative reduction in in-hospital mortality and a 3.3 percent absolute mortality reduction [14]. The system achieved an AUC of 0.97, identified 82 percent of sepsis cases in advance, and reached an 89 percent provider adoption rate. That last figure matters as much as any performance metric, since a tool clinicians don't use saves no one.

At Duke University Hospital, the Sepsis Watch system, built on 42 million data points from 42,000 patient encounters and assessing 86 variables every five minutes, predicted sepsis a median of five hours before clinical presentation and was associated with a 27 percent drop in sepsis-attributed deaths after deployment in November 2018 [15]. HCA Healthcare's Spot system, deployed across 185 hospitals, drew on 31 million annual patient care episodes as its training foundation and was estimated to have saved 8,000 lives over five years in conjunction with evidence-based interventions [16]. In my experience, the gap between a 97 AUC in the paper and what actually happens at 2 a.m. in a community hospital is where patients die. Algorithm performance is necessary; it's just not sufficient.

AI in Radiology and Pathology

The FDA had authorized more than 1,016 AI-enabled medical devices as of December 2024, with 96.4 percent cleared via the 510(k) pathway and radiology accounting for 76 percent of that total [17,18]. This concentration reflects where AI has found its most reliable clinical footing: pattern recognition in images.

For colonoscopy, a meta-analysis published in *Cancers* in 2025 found that computer-aided detection systems reduced the adenoma miss rate from 35.3

percent to 16.1 percent across tandem colonoscopy trials, a 54 percent relative reduction [19]. The GI Genius system from Medtronic, studied in a multicenter Danish quasi-randomized trial in 2025, increased the adenoma detection rate from 58.1 percent to 74.4 percent in the screening population [20]. Each 1 percent increase in adenoma detection rate corresponds to a 3 percent reduction in interval colorectal cancer risk [19].

In pathology, Paige AI's prostate cancer detection system became the first AI to receive FDA marketing authorization in pathology in 2021. A clinical validation study in which 18 pathologists reviewed 610 prostate biopsy whole-slide images prepared at 218 institutions found that AI assistance improved both sensitivity and specificity across all histologic grades and tumor sizes, with the stand-alone AI classification of biopsies as suspicious for cancer or benign aligning with the expert consensus reference standard [21]. What I tell residents is that AI in radiology doesn't replace the radiologist's judgment; it changes what the radiologist is looking at. The human eye moves from screening to confirmation, which is exactly where clinical expertise should be focused.

Large Language Models and Diagnostic Reasoning

Large language models have achieved impressive accuracy on standardized medical benchmarks. Google's Med-PaLM 2 achieved 86.5 percent accuracy on the MedQA USMLE-style benchmark [22]. GPT-4o reached 90.4 percent, with 92.7 percent accuracy on diagnostic reasoning questions [23]. These results represent genuine technical progress on standardized test questions. However, real clinical performance tells a different story. A 2025 meta-analysis in *npj Digital Medicine* covering 83 studies found that generative AI achieved only 52.1 percent accuracy when evaluated on actual clinical cases—substantially lower than the benchmark results and significantly inferior to specialist expert performance [24]. The discrepancy between 90+ percent benchmark accuracy and 52 percent clinical accuracy illustrates a fundamental challenge: test-taking ability doesn't guarantee real-world clinical judgment. No prospective randomized trial has demonstrated that LLM use improves patient outcomes.

Real-World Case Studies

The following cases illustrate how AI diagnostic tools perform under real clinical conditions, not in controlled research settings but in the daily workflow of practicing clinicians. Each case captures a different dimension of AI-assisted diagnosis: screening accuracy, integration with existing protocols, and the practical question of whether these tools change outcomes for actual patients.

Barbara and Mia

The gap between published performance metrics and patient outcomes can feel abstract until you meet Barbara.

Barbara is a woman who lives in Scotland. In 2024, she participated in the NHS Grampian breast screening pilot of an AI system called Mia, developed by Kheiron Medical Technologies. Two radiologists had reviewed her mammogram. Neither flagged anything of concern. Mia identified a 6 mm tumor, a tiny, early-stage cancer that was invisible to trained human eyes because of its size and characteristics. Because of that early detection, Barbara required only surgery and five days of radiotherapy. “It’s not a big C,” she told BBC News, “it’s a very little one.” Without Mia, her cancer almost certainly would have gone undetected until her next routine mammogram three years later, by which time the biological story would have been very different [25].

Barbara was not alone in that pilot. Across the 10,889 women whose mammograms Mia analyzed, the system identified 11 women with cancers that two radiologists had not flagged [25]. All were tiny tumors caught at their earliest, most treatable stage. Tumors under 15 millimeters at detection carry a 90 percent five-year survival rate. The tool also demonstrated the potential to cut results waiting time from 14 days to 3 days, a reduction with its own meaningful impact on patient anxiety and downstream care coordination.

The Barbara story captures something that performance statistics cannot: the experience of a person whose life was altered not by a physician’s genius or a miraculous new therapy but by a pattern-recognition system doing what pattern-recognition systems do best, finding a signal in noise

that a time-pressured human missed. It also lays out the moral arithmetic that healthcare leaders must confront: if a validated AI system can catch cancers that trained radiologists miss, the question of whether to deploy it is not just a quality improvement decision. It is an ethical one.

Spot

HCA Healthcare began developing Spot, the Sepsis Prediction and Optimization of Therapy system, because of a crisis that had no technological solution at the time: clinicians were missing the early signs of sepsis until patients were already critically ill. Sepsis is treatable. The problem isn't the treatment. It is the recognition. Spot was trained on data from 31 million annual patient care episodes across HCA's 185 hospitals, continuously monitoring vital signs, laboratory values, nursing assessments, and electronic health record data in real time. It was designed to detect what the human eye cannot: the subtle pattern of converging abnormalities that precede overt septic shock by hours. Over five years of deployment across 2.5 million patients, HCA reported that Spot, used in conjunction with evidence-based interventions, was associated with an estimated 8,000 lives saved [16]. The number is an institutional estimate, not a randomized trial result, and should be understood as such. But it comes from an organization with 31 million patient encounters annually and the data infrastructure to track outcomes at scale. It is the kind of figure that changes board conversations about AI investment.

TREWS

The most rigorously documented AI sepsis story is the one told by researchers at Johns Hopkins. TREWS was studied prospectively across 590,736 patients and five hospitals, with a companion paper in *Nature Medicine* examining the mechanisms of provider adoption alongside the mortality outcomes [14]. What the study found was not just that an AI system could identify sepsis. Many systems can do that with varying degrees of accuracy. It found that the way providers interacted with the system determined whether patients benefited. When a provider confirmed the TREWS alert and acted within three hours, the adjusted relative reduction in in-hospital mortality was 18.7 percent. When providers did not confirm the alert, the mortality benefit disappeared [14]. This finding

captures something fundamental about AI in clinical care: the tool is not the intervention. The tool is the prompt. The intervention remains the clinician.

The 89 percent provider adoption rate TREWS achieved, compared to 10 to 15 percent for legacy sepsis detection tools, was not accidental [14]. Bayesian Health, the company that commercialized TREWS, built the system's interface around one principle: make it easy for a clinician to act on what the algorithm is telling them and even easier to document why they are choosing not to. The result was a system that functioned as a genuine clinical partner rather than another source of alert fatigue. The lesson for health system leaders is direct: AI diagnostic tools that ignore workflow integration will fail regardless of their technical performance. The Google retinopathy screening study in Thailand showed that a system with 90 percent accuracy in a research setting can fail completely when deployed in a clinical environment where nurses lack the training to handle edge cases and patients lack reliable access to follow-up care. Performance in the lab and performance in the clinic are different things, measured in different units.

Risks and Limitations

The evidence for AI predictive diagnostics is real, and it is growing. A responsible reading of that evidence requires equal attention to what the research reveals about its limits, not as a reason for inaction, but as a roadmap for deployment that is honest about trade-offs.

Overdiagnosis and Lead-Time Bias

The most fundamental challenge in screening, human or AI-assisted, is distinguishing detection from improved outcomes. An analysis of the National Lung Screening Trial found that approximately 20 percent of cancers detected by low-dose CT represented overdiagnosis: cancers that would never have become clinically apparent [26]. As Adamson and Welch observed in *NEJM*, AI tools may detect more consistently than human readers, "but they won't necessarily be closer to the truth, algorithms may not be any better than humans at determining which tumors are destined to cause symptoms or death" [27].

The MASAI trial's 12 percent reduction in interval cancers is a meaningful outcome proxy. Long-term mortality data, the standard that separates screening advances from screening mirages, is not yet available [5].

False Positives, Cascades, and Disparities

False positives in AI screening are not simply incorrect results that clinicians discard. They are clinical events with psychological, financial, and behavioral consequences. A 2024 study in the *Journal of Epidemiology and Community Health* found that women with false-positive screening results used GP services 1.1 to 1.2 times more than screening-naive women for years afterward and that false positives were associated with 10 to 154 percent higher rates of subsequent breast imaging [28]. The psychosocial consequences of a false positive can persist for up to 14 years. Perhaps most counterintuitively, a UC Davis study found that after a false positive, only 61 percent of women returned for subsequent mammography screening, versus 77 percent following a negative result [29]. The psychological impact of false positives is complex: affected women tend to seek more immediate medical attention and additional imaging, but they are paradoxically less likely to return for future routine screening, the very behavior that early detection depends on.

Racial disparities in AI false positive rates represent an equity problem embedded inside a quality improvement tool. A 2024 study in *Radiology* examined 4,855 true-negative screening digital breast tomosynthesis exams and found that Black patients had 1.5 times higher odds of false-positive case scores compared to white patients and that patients aged 61 to 70 had 3.5 times higher odds of false-positive risk scores compared to patients aged 51 to 60 [30]. These are not statistical abstractions. They are the differential burden of unnecessary biopsies, unnecessary anxiety, and unnecessary downstream procedures disproportionately carried by populations that already bear a greater burden of diagnostic delay.

Generalizability Failures

A Johns Hopkins review of 86 deep-learning radiology algorithms found that 81 percent showed decreased accuracy on external datasets compared to internal validation and that nearly 24 percent experienced a substantial drop, defined as a drop of 0.10 AUC or more, when tested outside the

institution where they were developed [31]. The Epic Sepsis Model, deployed at hundreds of U.S. hospitals based on internal performance documentation showing an AUC of 0.76 to 0.83, was externally validated at Michigan Medicine and found to have an AUC of 0.63, a number that should have prevented widespread deployment [32]. At a decision threshold of 6, the model missed 67 percent of sepsis patients while generating alerts in 18 percent of all hospitalizations, with a positive predictive value of just 12 percent [32]. The result was a system that simultaneously flooded clinicians with nonactionable alerts and failed to identify most patients who needed intervention. It is the cautionary archetype for what happens when AI tools are deployed without rigorous external validation.

Black Box Trust, Automation Bias, and Legal Exposure

The opacity of deep learning systems creates a specific clinical danger. A study examining nursing students and licensed nurses using AI early warning tools found that when AI predictions were misleading, nurses' clinical decision-making error rates increased by 96 to 120 percent compared to their performance without AI assistance [33]. The problem was not that clinicians followed incorrect advice. It was that they couldn't reliably detect when the AI was wrong. This is automation bias: the tendency to defer to a machine's output without maintaining the independent judgment required to catch machine errors.

In my experience, this is the failure mode that clinical training most consistently underestimates. We teach physicians to use tools. We rarely teach them to suspect the tools.

The legal consequences of this dynamic are already materializing. A mock jury study involving more than 1,300 participants found that when an AI system correctly identified a brain bleed and a radiologist missed it, jurors sided with the plaintiff 74 percent of the time, compared to 56 percent when no AI was described [34]. Radiologists are being held to a higher standard of accountability not because the law has changed, but because the public expectation of what a reasonable physician should do in an AI-assisted environment is changing. Malpractice claims involving AI diagnostic tools increased 14 percent between 2022 and 2024 [35].

Cost-Effectiveness Questions

The economics of AI diagnostics are not uniformly favorable. A 2025 microsimulation analysis of AI-assisted breast cancer screening found an incremental cost-effectiveness ratio of $303,279 per quality-adjusted life year, a figure that was above standard willingness-to-pay thresholds in 98 percent of simulations [36]. A quality-adjusted life year (QALY) is a standard health economics measure combining both the quantity and quality of life generated by a healthcare intervention. Cost-effectiveness thresholds in the U.S. typically consider interventions below $50,000 to $100,000 per QALY as cost-effective. This doesn't mean AI mammography is without value. It means that at current pricing, the economic case has not been made. By contrast, Mayo Clinic's AI-ECG screening for left ventricular dysfunction achieves a cost-effectiveness ratio of $1,651 per QALY in outpatient settings, a threshold that no rational health economist would contest [37].

The lesson is not uniformity but context: AI diagnostic tools must be evaluated individually, in specific clinical settings, against specific comparators, with specific patient populations. The blanket assertion that AI is cost-effective or cost-ineffective is equally off-base.

Evidence Gap on Patient Outcomes

A systematic review of 39 randomized controlled trials found that only 70 percent of AI trials with positive outcomes used clinically relevant endpoints rather than diagnostic performance metrics [38]. Detection rates are not mortality rates. The evidence base for AI diagnostics is maturing, but the fundamental question, do patients live longer, remains incompletely answered.

Shortcut learning inflates reported AI accuracy by an average of 20 percent across training datasets involving 207,487 patients, 5 modalities, and 13 clinical datasets [39]. Models have learned to detect chest catheters as a proxy for pneumonia, and scanner-type as a proxy for disease, spurious correlations that collapse when the model encounters a different institution, patient population, or imaging protocol.

Regulatory and Legal Landscape

The regulatory environment for AI diagnostics in the United States has grown more complex, more sophisticated, and in some ways more ambiguous than at any prior point.

FDA Authorization and Pathways

As of December 2024, the FDA had authorized 1,016 AI and machine learning-enabled medical devices, with 96.4 percent cleared via the 510(k) pathway based on substantial equivalence to a predicate device [17]. Radiology accounts for 76 percent of all authorized AI devices. Cardiovascular tools account for 9.8 percent [17]. The de novo pathway, used for truly novel devices without a predicate, such as IDx-DR and Paige Prostate, has been used in approximately 3 percent of cases. Premarket approval, the highest-risk pathway requiring clinical evidence of safety and effectiveness, accounts for less than 1 percent.

In December 2024, the FDA finalized guidance on Predetermined Change Control Plans, which allow manufacturers to specify planned AI algorithm modifications in their initial submission and implement those changes without a new premarket review, a meaningful step toward adaptive lifecycle management [40]. In January 2025, the FDA issued draft guidance on AI-enabled device lifecycle management, proposing full requirements for ongoing monitoring, performance evaluation, and drift detection [40].

The 21st Century Cures Act created an exemption from device regulation for clinical decision support software that meets four specific criteria, including the requirement that a healthcare professional be able to independently review the basis for any recommendation. Software that analyzes imaging data, ECG signals, or predicts time-critical events does not qualify for this exemption and remains subject to FDA oversight [41].

The Laboratory-Developed Test (LDT) Regulation Saga

One of the most consequential, and turbulent, regulatory episodes in recent healthcare history unfolded over the past two years. In May 2024, the FDA issued a final rule that would have subjected approximately 80,000 laboratory-developed tests, including those incorporating AI or machine

learning algorithms, to premarket review requirements for the first time [42]. The rule was vacated in March 2025 by a federal district court, which held that Congress had already addressed LDT regulation through the Clinical Laboratory Improvement Amendments (CLIA) framework administered by the Centers for Medicare & Medicaid Services (CMS), and that the FDA's assertion of jurisdiction "defies bedrock principles of statutory, common sense, and longstanding industry practice" [42]. In September 2025, the FDA issued a rescission rule restoring enforcement discretion for LDTs to its pre-2024 baseline [42]. The practical consequence: AI-enhanced laboratory tests developed and used within a single certified laboratory remain, for now, outside the premarket review system.

Reimbursement: The Coverage Frontier

Authorization is necessary but not sufficient for clinical deployment. Reimbursement determines whether authorized tools can be economically sustained. As of January 2026, there are 26 CPT codes for clinical AI solutions, with only 3 having achieved Category I status, the highest level of permanence and reimbursement [43]. The Transitional Coverage for Emerging Technologies pathway, finalized in 2024, creates an expedited route to Medicare coverage for certain FDA Breakthrough Device–designated tools, with CMS aiming to finalize national coverage determinations within six months of FDA authorization [44].

The EU AI Act, which entered into force in August 2024 with high-risk provisions applicable from 2025 and 2026, classifies most healthcare AI systems, including diagnostic software requiring notified-body review under the Medical Device Regulation, as high-risk AI, triggering requirements for data governance, transparency, human oversight, robustness testing, and quality management [45].

Standard of Care Evolution

The legal standard for AI in clinical practice is shifting under the pressure of the technology's capabilities. A 14 percent increase in malpractice claims involving AI diagnostic tools between 2022 and 2024 reflects not a system dysfunction but an expectation adjustment [35]. The Federation of State Medical Boards recommended in 2024 that clinicians, not AI developers,

bear liability for AI-related diagnostic errors [46]. The AMA's November 2024 AI Principles document raised explicitly the question that will define the next decade of healthcare law: "As AI continues to evolve, there may ultimately be questions regarding liability when physicians fail to use AI and rely only on their professional judgment" [47].

The American College of Radiology (ACR) launched its Assess-AI registry in November 2024, the world's first AI quality registry, which monitors ongoing performance of imaging AI algorithms in real-world clinical settings and provides facilities with analytics on local performance against national benchmarks [48]. The AMA released its Governance for Augmented Intelligence toolkit in August 2025, providing an eight-step framework for health systems to implement, evaluate, and scale AI, from defining permitted use cases to establishing training requirements and accountability structures [47].

Ethical Implications

The ethical terrain of predictive AI diagnostics requires navigating two principles that, in clinical practice, frequently pull against each other: the duty to detect and the right not to know.

Right Not to Know and the Duty to Detect

Bioethics has long recognized the right not to know as a patient's qualified claim to decline predictive health information about themselves. This right is grounded in personal autonomy and is recognized in international frameworks including the UNESCO Universal Declaration on the Human Genome, the Oviedo Convention, and the CIOMS Guidelines [49]. But it is a qualified right, not absolute. It must be balanced against the duty of care when findings are clinically actionable, the duty to prevent harm when risk is high and intervention effective, and, in genetic contexts, the duty to third parties who share biological inheritance.

Predictive AI creates scenarios that were practically impossible in the era of reactive medicine: a patient who has not sought diagnosis may nonetheless be identified by an AI screening algorithm as carrying a 40 percent risk of developing pancreatic cancer in the next three years. What is the clinician's

obligation in that moment? What is the patient's right? These questions have no clean answers. What institutions can do is establish explicit frameworks for how they will be navigated, before the AI tools are deployed, not after the first difficult case arrives.

Incidental Findings at Scale

Between 15 and 30 percent of all diagnostic imaging already contains at least one incidental finding [50]. AI systems applied to imaging will systematically increase that proportion, because they are designed to detect subtler signals than radiologists can reliably perceive. The clinical and ethical challenge is distinguishing actionable incidental findings from noise that, if acted upon, causes net harm. Hospitals already spend an estimated $43 million annually settling lawsuits for missed follow-up on incidental lung findings alone [51]. AI-scale detection will intensify this challenge without automatically providing the clinical resources to address it.

The Equity Imperative

Equity in access to AI diagnostics is not a secondary concern. It is the central test of whether AI diagnostic tools deliver on their potential or amplify the disparities they were presumed to reduce. A 2019 study in *Science* found that a widely deployed health system algorithm exhibited systematic racial bias, with Black patients having 26.3 percent more chronic illnesses than white patients at the same risk score, because the algorithm used healthcare costs as a proxy for health needs and systemic barriers to access mean that Black patients generate lower costs despite equal or greater illness burden [52]. A meaningful share of office-based practices, particularly small, solo, rural, and safety-net practices, continue to lag in the health IT infrastructure that most AI systems require, disproportionately affecting rural, uninsured, and underserved populations [53].

A predictive AI diagnostic tool that performs well in a well-resourced academic medical center and poorly in an under-resourced community clinic is not a solution to the diagnostic error crisis. It transfers diagnostic advantage from patients who already have less of it to patients who already have more.

Automation Bias and Clinician Training

The automation bias problem is real and measurable. When AI predictions are correct, clinicians perform significantly better. When AI predictions are misleading, clinicians perform significantly worse than they would have without the AI. In one study, their error rates were 96 to 120 percent higher [33]. This asymmetry has a clear policy implication: deploying AI diagnostic tools without structured training on their failure modes isn't simply inefficient. It is potentially dangerous. Clinicians need to know not only how to use these tools but specifically when to distrust them.

The WHO's six ethical principles for AI in health, protecting human autonomy, promoting well-being and the public interest, ensuring transparency and explainability, building accountability, ensuring inclusiveness and equity, and promoting sustainability, provide a useful organizing framework [54]. Applied to predictive diagnostics, they translate into concrete institutional obligations: AI tools must be explainable to the clinicians using them, must be evaluated for performance across demographic subgroups before deployment, must include human oversight mechanisms that are meaningful rather than nominal, and must be subject to ongoing monitoring as patient populations and clinical workflows evolve.

The tension between early detection and patient anxiety is real. Patients who receive a positive result from a multicancer early detection test and are then told their cancer diagnosis could not be confirmed may experience lasting psychological harm. A 2025 study in *npj Digital Medicine* found that when AI mammography scores were disclosed alongside normal radiologist reports, 91.5 percent of patients expressed concern about AI misdiagnosis, and the desire for additional follow-up and specialist consultation rose significantly even when both the radiologist and the AI agreed the scan was normal [55]. Transparency about AI results is ethically required. So is careful attention to how that transparency is communicated.

Leadership Implementation Framework

For healthcare executives, the question is not whether AI predictive diagnostics work. The evidence reviewed in this chapter shows that in specific, validated contexts, they do. The question is how to build the

institutional infrastructure to deploy them responsibly, monitor them rigorously, and update or retire them when the evidence warrants.

Governance Models from Leading Institutions

The institutions that have moved most effectively into AI diagnostic deployment share a governance model built on three elements: multidisciplinary oversight, continuous performance monitoring, and explicit clinical workflow integration.

Kaiser Permanente Kaiser Permanente's predictive early-warning system, which the organization estimates reduces hospital mortality by approximately 500 deaths per year and readmissions by 10 percent, is governed by seven principles of responsible AI that include a clinician oversight requirement: all AI-generated content must be reviewed by a clinician before being acted upon [56]. Patients are permitted to opt out of AI-assisted scribing. The governance structure existed before the tools were deployed, not after.

Mayo Clinic Mayo Clinic has invested in more than 200 AI projects enterprise-wide, created a Department of AI and Informatics, and formed a 10-year strategic partnership with Google Cloud for co-developed AI solutions [56]. Its AI-ECG heart failure screening program, proven in randomized trials to significantly increase early detection of low ejection fraction at a cost of $1,651 per QALY in outpatient settings [37], represents the kind of AI investment that a health system can defend to a board of trustees, a quality committee, and a malpractice insurer simultaneously.

Duke Health Duke Health coined the concept of "algorithm vigilance," continuous post-deployment monitoring of AI system performance across patient subgroups, with explicit bias detection built into the monitoring protocol [56]. Sepsis Watch's post-deployment mortality reduction of 27 percent at Duke University Hospital was not achieved by deploying an algorithm and waiting for outcomes to improve. It was achieved by building a centralized team of rapid-response nurses to act on the algorithm's output, creating a new clinical role that the algorithm made possible and necessary [15].

Cleveland Clinic Cleveland Clinic formed a multidisciplinary AI oversight team including clinicians, data scientists, ethicists, and information technology staff before deploying AI tools at scale [56].

UCSF Health UCSF Health hired an inaugural chief health AI officer and built a secure internal LLM cloud architecture ensuring that all prompts and patient data remain within the institution's governance boundary [56].

I've found that the common thread in these governance models isn't budget or institutional size. It's the decision to treat AI like a new medication: test it, monitor it, and be willing to pull it when it stops working as intended.

The ACR Assess-AI Registry and AMA Governance Toolkit

Two professional society resources deserve specific attention in any institutional AI governance framework. The ACR's Assess-AI registry, launched in November 2024, allows radiology facilities to monitor the real-world performance of AI imaging algorithms against national benchmarks, providing the kind of continuous feedback loop that static validation studies cannot [48]. For any institution deploying AI diagnostic imaging tools, enrollment in Assess-AI is an emerging best practice.

The AMA's Governance for Augmented Intelligence toolkit provides an eight-step implementation framework [47]:

- Define AI terms and categories
- Articulate AI risks
- Establish permitted and prohibited use cases
- Set governance, evaluation, and approval processes
- Define accountability and oversight structures
- Establish data retention policies
- Set transparency guidelines for disclosure to providers and patients
- Define training requirements

The CHAI framework, developed by the 3,000-member Coalition for Health AI organization, extends this to a voluntary certification model that the Joint Commission endorsed in 2025 through the first guidance on Responsible Use of AI in Healthcare [56].

Continuous Monitoring and Drift Detection

AI systems that perform well at deployment may degrade over time as patient populations shift, clinical practices change, or the underlying data distribution drifts from the training distribution. The FDA formally acknowledges that AI performance "can be influenced by changes in clinical practice, patient demographics, data inputs, and health care infrastructure," and that current retrospective testing is insufficient to predict real-world behavior [57]. Institutions must build post-deployment monitoring into the original implementation plan, not treat it as an afterthought.

Scheduled performance reviews at six-month intervals, stratified by demographic subgroup, should be standard. Tools that show evidence of performance degradation should be flagged for re-validation or retirement. The standard for AI diagnostic tools should be at least as rigorous as the standard applied to clinical laboratory equipment. The CLIA model of ongoing proficiency testing and quality control is not a ceiling but a floor.

Future Outlook

The AI diagnostic tools described in this chapter, trained on individual data modalities, validated in specific clinical contexts, deployed as stand-alone applications, are in historical terms the first generation. What comes next is qualitatively different.

Foundation Models for Diagnostics

Foundation models, large AI systems pre-trained on massive, diverse datasets that can be adapted to numerous downstream tasks with minimal additional training, are already reshaping AI in medicine. BiomedGPT and MedSAM represent early demonstrations of generalized performance across multiple imaging modalities without task-specific architectures [58].

Aidoc's Care platform, which received the first-ever FDA Breakthrough Device Designation for a multitriage AI system covering a wide array of life-threatening conditions in a single CT workflow, represents the first clinical-grade foundation model with FDA-cleared solutions [59]. These systems don't merely analyze a single finding on a single scan. They analyze a patient's entire imaging history simultaneously, looking for patterns across studies, time points, and anatomical regions.

Federated Learning

Federated learning, a distributed approach where AI models are trained locally at individual institutions with only model parameters, not raw patient data, shared to a central server, addresses the most fundamental obstacle to building generalizable medical AI: the fact that most clinical data can't legally or ethically leave the institution where it was generated. The Eye2Gene project, which identifies genetic causes of inherited retinal disease from retinal scans with 86 percent accuracy for its top five candidates, was trained across institutions in the United Kingdom, Germany, Australia, Brazil, Singapore, France, and Japan, with data never leaving any originating hospital [60]. The global federated learning in healthcare market is projected to grow from $35.67 million in 2025 to $141.01 million by approximately 2034 [60].

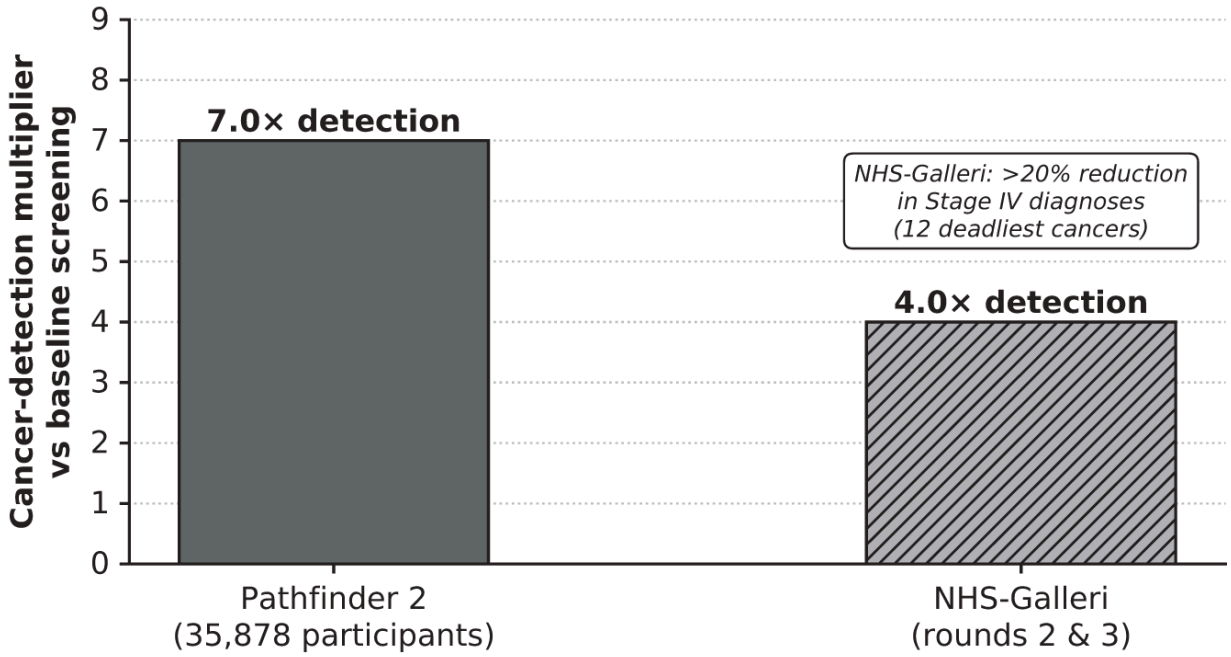
Multimodal AI and the Integration of Evidence

The next frontier in predictive diagnostics is not imaging AI or genomic AI or EHR AI. It is the integration of all three. Multimodal AI systems that combine CT imaging with electronic health records, laboratory results, genomic variants, and continuous wearable data will identify patients at risk of specific diseases years before either any single modality or any human clinician could. The NIH All of Us program, building a diverse database of one million or more U.S. participants with longitudinal health data, represents the kind of training resource that will make these systems feasible at population scale.

The multicancer early detection field is moving toward population-level screening. Pathfinder 2 demonstrated a sevenfold increase in cancer detection in a 35,878-participant study across 32 healthcare systems [9]. The NHS-Galleri trial, despite not meeting its primary endpoint,

demonstrated a fourfold higher cancer detection rate and a greater than 20 percent reduction in Stage IV diagnoses for the 12 deadliest cancer types in the second and third screening rounds [10]. The question is no longer whether blood-based multicancer detection is real. The question is how to identify which patients are most likely to benefit, how to manage the downstream diagnostic workload that positive results generate, and how to make these tests accessible across economic and geographic lines.

The vision is medicine that is predictive rather than reactive, meeting disease before it becomes catastrophic, extending the window of effective intervention from months to years. Parts of that vision are already real: in a breast screening clinic in Scotland, in the sepsis wards of Johns Hopkins, in the cardiology offices where a 10-second ECG predicts a stroke years before it occurs. The task for healthcare leaders is not to wait for perfection. It is to build the governance structures, monitoring systems, and clinical culture that allow these tools to deliver on their promise while protecting patients from the risks that come with them (see Figure 4.3).



Source: Pathfinder 2 study (32 health systems); NHS-Galleri trial (UK).

Figure 4.3 AI versus human detection performance across cancer types

Key Takeaways

Diagnostic errors kill or permanently disable an estimated 795,000 Americans annually. AI predictive diagnostics represent the most significant opportunity to reduce this toll in the history of modern medicine, but only if deployed with the rigor the evidence demands.

The MASAI trial is the foundational evidence for AI mammography. The trial demonstrated 29 percent more cancers detected, 12 percent fewer interval cancers, and 44 percent reduced radiologist workload without increasing false positive rates. For health system leaders considering AI mammography, this is where the decision begins.

TREWS achieved an 18.7 percent relative reduction in sepsis mortality across 590,736 patients. The system also achieved 89 percent provider adoption, the highest ever reported for a clinical decision support tool. Workflow design matters as much as algorithm performance. An AI tool that clinicians don’t use saves no one.

Eighty-one percent of deep learning radiology algorithms showed decreased accuracy on external datasets. Require external validation data from comparable institutions before deployment. Internal performance metrics are not sufficient.

Automation bias can make clinicians perform worse than if they had no AI at all. When AI predictions were misleading, nurses performed 96 to 120 percent worse than without AI. Structured training on AI failure modes is a patient safety requirement, not optional professional development.

Demographic disparities in AI diagnostic performance are documented and measurable. Black patients had 1.5 times higher odds of false-positive AI scores in one breast screening study. Demographic performance stratification must be a prerequisite for institutional AI approval, not an afterthought.

Reimbursement gaps are blocking clinical deployment. Only 3 of 26 AI CPT codes had achieved Category I reimbursement status as of January 2026. Health system leaders must engage the reimbursement

policy process. Authorization without coverage is not clinical deployment.

Economic analysis must precede deployment decisions. AI-ECG for heart failure screening costs $1,651 per QALY in outpatient settings; AI-assisted breast cancer screening has been modeled at $303,279 per QALY at current pricing. Context-specific economic analysis is not optional, it determines whether a tool that works clinically can be sustained financially.

Governance built before deployment is now a legal requirement. The institutions achieving the best AI outcomes, Kaiser, Mayo, Duke, Cleveland Clinic, UCSF Health, built governance structures, monitoring protocols, and accountability systems before deploying tools at scale. A 14 percent increase in malpractice claims involving AI diagnostic tools between 2022 and 2024 confirms that governance is now a legal requirement, not just a quality improvement choice.

References

1. Newman-Toker DE, Nassery N, Schaffer AC, Yu-Moe CW, Clemens GD, Wang Z, et al. Burden of serious harms from diagnostic error in the USA. BMJ Qual Saf. 2024; 33(2):109–120. doi:10.1136/bmjqs-2021-014130.

2. Singh H, Meyer AN, Thomas EJ. Frequency of diagnostic errors in outpatient care. BMJ Qual Saf. 2014; 23(9):727–731. Available from: https://qualitysafety.bmj.com/content/23/9/727.

3. Agency for Healthcare Research and Quality. Diagnostic errors in the emergency department. Rockville (MD): AHRQ; 2022. Available from: https://www.ncbi.nlm.nih.gov/books/NBK588113.

4. Adverse diagnostic events in hospitalised patients. BMJ Qual Saf. 2024. Available from: https://bmjgroup.com/harmful-diagnostic-errors-may-occur-in-1-in-every-14-general-medical-hospital-patients.

5. Lång K, Josefsson V, Larsson AM, Larsson S, Högberg C, Sartor H, et al. Artificial intelligence-supported screen reading versus standard

double reading in the Mammography Screening with Artificial Intelligence trial (MASAI): a clinical safety analysis of a randomised, controlled, non-inferiority, single-blinded, screening accuracy study. Lancet Oncol. 2023; 24(8):936–944. doi:10.1016/S1470-2045(23)00298-X.

6. Eisemann N, Bunk S, Mukama T, Baltus H, Elsner SA, Gomille T, et al. Nationwide real-world implementation of AI for cancer detection in population-based mammography screening. Nat Med. 2025; 31(3):917–924. doi:10.1038/s41591-024-03408-6.

7. Mikhael PG, Wohlwend J, Yala A, Karstens L, Xiang J, Takigami AK, et al. Sybil: a validated deep learning model to predict future lung cancer risk from a single low-dose chest computed tomography. J Clin Oncol. 2023; 41(12):2191–2200. doi:10.1200/JCO.22.01345.

8. Sybil AI Model Demonstrates Strong Performance in Predicting Lung Cancer Risk with Minimal Bias. ILCN (IASLC) [Internet]. 2025 Sep 7 [cited 2026 Jul 5]. Available from: https://www.ilcn.org/sybil-ai-model-demonstrates-strong-performance-in-predicting-lung-cancer-risk-with-minimal-bias.

9. GRAIL. GRAIL PATHFINDER 2 Results Show Galleri Multi-Cancer Early Detection Blood Test Increased Cancer Detection More Than Seven-Fold When Added to USPSTF A and B Recommended Screenings [press release]. Menlo Park (CA): GRAIL; 2025 Oct 17. Available from: https://grail.com/press-releases/grail-pathfinder-2-results-show-galleri-multi-cancer-early-detection-blood-test-increased-cancer-detection-more-than-seven-fold-when-added-to-uspstf-a-and-b-recommended-screenings.

10. GRAIL. NHS-Galleri trial top-line results. 2026. Available from: https://grail.com/press-releases/landmark-nhs-galleri-trial-demonstrates-a-substantial-reduction-in-stage-iv-cancer-diagnoses-increased-stage-i-and-ii-detection-of-deadly-cancers-and-four-fold-higher-cancer-detection-rate.

11. Attia ZI, Noseworthy PA, Lopez-Jimenez F, Asirvatham SJ, Deshmukh AJ, Gersh BJ, et al. An artificial intelligence–enabled ECG algorithm for

the identification of patients with atrial fibrillation during sinus rhythm. Lancet. 2019; 394(10201):861–867.

12. Shahid S, Iqbal M, Saeed H, Hira S, Batool A, Khalid S, et al. Diagnostic accuracy of Apple Watch electrocardiogram for atrial fibrillation: a systematic review and meta-analysis. JACC Adv. 2025; 4(2):101538.

13. van Steijn NJ, Blommestijn IS, Blok S, Pepplinkhuizen S, Somsen GA, Knops RE, et al. Enhanced detection and prompt diagnosis of atrial fibrillation using Apple Watch. J Am Coll Cardiol. 2026; 87(14):1714–1728. doi:10.1016/j.jacc.2025.11.032.

14. Adams R, Henry KE, Sridharan A, Soleimani H, Zhan A, Rawat N, et al. Prospective, multi-site study of patient outcomes after implementation of the TREWS machine learning-based early warning system for sepsis. Nat Med. 2022; 28(7):1455–1460. doi:10.1038/s41591-022-01894-0.

15. Duke Institute for Health Innovation. Sepsis Watch. Durham (NC): Duke University; 2024. Available from: https://dihi.org/project/sepsiswatch.

16. HCA Healthcare. Algorithm-driven sepsis detection improves outcomes. 2019. Available from: https://investor.hcahealthcare.com/news/news-details/2019/HCA-Healthcare-Using-Algorithm-Driven-Technology-to-Detect-Sepsis-Early-and-Help-Save-8000-Lives/default.aspx.

17. Mehta V, Komanduri A, Bhadouriya RS, Mehta V, Johnson MD, et al. Evaluating transparency in AI/ML model characteristics for FDA-reviewed medical devices. NPJ Digit Med. 2025; 8:673. doi:10.1038/s41746-025-02052-9.

18. Almarie B, Gonzalez-Gonzalez LF, dos Santos Barbosa LA, Lutz A, Grosse U, Fregni F, et al. Machine learning-enabled medical devices authorized by the US Food and Drug Administration in 2024: regulatory characteristics, predicate lineage, and transparency reporting. Biomedicines. 2025; 13(12):3005. doi:10.3390/biomedicines13123005.

19. Spadaccini M, Menini M, Massimi D, Rizkala T, De Sire R, Alfarone L, et al. AI and polyp detection during colonoscopy. Cancers (Basel). 2025; 17(5):797. doi:10.3390/cancers17050797.

20. Lagström RMB, Bräuner KB, Bielik J, Rosen AW, Crone JG, Gögenur I, et al. Improvement in adenoma detection rate by artificial intelligence-assisted colonoscopy: multicenter quasi-randomized controlled trial. Endosc Int Open. 2025; 13:a25215169. doi:10.1055/a-2521-5169.

21. Raciti P, Sue J, Retamero JA, Ceballos R, Godrich R, Kunz JD, et al. Clinical validation of artificial intelligence-augmented pathology diagnosis demonstrates significant gains in diagnostic accuracy in prostate cancer detection. Arch Pathol Lab Med. 2023; 147(10):1178–1185. doi:10.5858/arpa.2022-0066-OA.

22. Singhal K, Azizi S, Tu T, Mahdavi SS, Wei J, Chung HW, et al. Large language models encode clinical knowledge. Nature. 2023; 620(7972):172–180. doi:10.1038/s41586-023-06291-2.

23. Bicknell BT, Butler D, Whalen S, Ricks J, Dixon CL, Clark AB, et al. ChatGPT-4 Omni performance in USMLE disciplines and clinical skills: comparative analysis. JMIR Med Educ. 2024; 10:e63430. doi:10.2196/63430.

24. Takita H, Kabata D, Walston SL, Tatekawa H, Saito K, Tsujimoto Y, et al. A systematic review and meta-analysis of diagnostic performance comparison between generative AI and physicians. NPJ Digit Med. 2025; 8:175. doi:10.1038/s41746-025-01543-z.

25. Kleinman Z. NHS pilots AI breast cancer screening tool. BBC News. 2024. Available from: https://www.bbc.com/news/technology-68607059.

26. Patz EF Jr, Pinsky P, Gatsonis C, Sicks JD, Kramer BS, Tammemägi MC, et al. Overdiagnosis in low-dose computed tomography screening for lung cancer. JAMA Intern Med. 2014; 174(2):269–274. doi:10.1001/jamainternmed.2013.12738.

27. Adamson AS, Welch HG. Machine learning and the cancer-diagnosis problem — no gold standard. N Engl J Med. 2019; 381(24):2285–2287.

doi:10.1056/NEJMp1907407.

28. Gram EG, Siersma V, Nicolaisdóttir DR, Brandt Brodersen J. Downstream healthcare use following breast cancer screening: a register-based cohort study. J Epidemiol Community Health. 2025; 79(4):242–248. doi:10.1136/jech-2024-222818.

29. Nordic Global. Improving follow-up after false-positive screening. 2024. Available from: https://www.nordicglobal.com/blog/from-anxiety-to-action-human-centered-solutions-to-increase-follow-up-after-false-positives.

30. Nguyen DL, Ren Y, Jones TM, Thomas SM, Lo JY, Grimm LJ. Patient characteristics impact performance of AI algorithm in interpreting negative screening digital breast tomosynthesis studies. Radiology. 2024; 311(2):e232286. doi:10.1148/radiol.232286.

31. Yu AC, Mohajer B, Eng J. External validation of deep learning algorithms for radiologic diagnosis: a systematic review. Radiol Artif Intell. 2022; 4(3):e210064. doi:10.1148/ryai.210064.

32. Wong A, Otles E, Donnelly JP, Krumm A, McCullough J, DeTroyer-Cooley O, et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients (the Epic Sepsis Model). JAMA Intern Med. 2021; 181(8):1065–1070. doi:10.1001/jamainternmed.2021.2626.

33. How AI can degrade human performance in high-stakes settings. AI Frontiers. 2025. Available from: https://ai-frontiers.org/articles/how-ai-can-degrade-human-performance-in-high-stakes-settings.

34. Brown University. AI liability in radiology: mock jury study. NEJM AI. 2025. Available from: https://medical.brown.edu/news/2025-07-28/radiology-artificial-intelligence-malpractice-study.

35. Broderick B. Medical malpractice trends in the AI era. 2025. Available from: https://www.brandonjbroderick.com/medical-malpractice-2025-how-ai-healthcare-changing-lawsuits.

36. Andersen M, Kunst N, Richman IB, et al. Long-term outcomes and cost-effectiveness of artificial intelligence for breast cancer screening: a modeling study. Value Health. 2026; 29(2):179–188. PMID: 40983310.

37. Thao V, Zhu Y, Tseng AS, Inselman JW, Borah B, et al. Cost-effectiveness of artificial intelligence-enabled electrocardiograms for early detection of low ejection fraction: a secondary analysis of the electrocardiogram artificial intelligence-guided screening for low ejection fraction trial. Mayo Clin Proc Digit Health. 2024; 2(4):620–631. doi:10.1016/j.mcpdig.2024.10.001.

38. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. BMJ. 2022; 377:e070904. doi:10.1136/bmj-2022-070904.

39. Ong Ly C, Unnikrishnan B, Tadic T, Patel T, Duhamel J, Kandel S, et al. Shortcut learning in medical AI hinges on high-quality data. NPJ Digit Med. 2024; 7:124. doi:10.1038/s41746-024-01118-4. PMID 38744921.

40. US Food and Drug Administration. Artificial intelligence software as a medical device (SaMD). 2025. Available from: https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-software-medical-device.

41. US Food and Drug Administration. Clinical decision support software guidance. 2026. Available from: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software.

42. US Food and Drug Administration7. Laboratory developed tests (LDTs). 2025. Available from: https://www.fda.gov/medical-devices/in-vitro-diagnostics/laboratory-developed-tests.

43. Bipartisan Policy Center. Paying for AI in US health care. 2026. Available from: https://bipartisanpolicy.org/issue-brief/paying-for-ai-in-u-s-health-care.

44. Reuter E. More than 1,100 devices have received the FDA's breakthrough designation. MedTech Dive [Internet]. 2025 Aug 21 [cited

2026 Jul 5]. Available from: https://www.medtechdive.com/news/fda-breakthrough-device-update-2025/758292.

45. Sobande S. EU AIA and its interplay with the EU MDR/IVDR [Internet]. Emergo by UL; 2025 Mar 20 [cited 2026 Jul 5]. Available from: https://www.emergobyul.com/resources/eu-aia-and-its-interplay-eu-mdrivdr.

46. Catherman C. Are doctors liable when AI makes a mistake? Malpractice experts weigh in. Healthcare Brew. 2025 Apr 1. Available from: https://www.healthcare-brew.com/stories/2025/04/01/doctors-liable-ai-mistake-malpractice.

47. American Medical Association. AI principles and governance toolkit. 2025. Available from: https://www.ama-assn.org/system/files/ama-ai-principles.pdf.

48. American College of Radiology. Assess-AI registry. 2024. Available from: https://www.acr.org/News-and-Publications/Media-Center/2024/American-College-of-Radiology-Launches-Landmark-Artificial-Intelligence-Quality-Registry.

49. Timelex. Right not to know in health research. 2026. Available from: https://www.timelex.eu/en/blog/right-not-know-and-incidental-findings-health-research-balancing-autonomy-and-duty-inform.

50. Dudum R, Asch FM. “Incidentalomas” in the age of artificial intelligence. J Gen Intern Med. 2023; 38(13):2855–2856. doi:10.1007/s11606-023-08325-x.

51. Northwestern Medicine. Using artificial intelligence to solve one of health care’s most enduring problems [Internet]. Chicago (IL): Northwestern Memorial Hospital; 2022 Mar 16 [cited 2026 Jul 5]. Available from: https://news.nm.org/using-artificial-intelligence-to-solve-one-of-health-cares-most-enduring-problems.

52. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019; 366(6464):447-453. doi:10.1126/science.aax2342.

53. Hadjiat Y. Healthcare inequity and digital health—a bridge for the divide, or further erosion of the chasm? PLOS Digit Health. 2023; 2(6):e0000268. doi:10.1371/journal.pdig.0000268.

54. World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: World Health Organization; 2021. Available from: https://www.who.int/publications/i/item/9789240029200.

55. Khullar D, Casalino LP, Qian Y, Lu Y, Krumholz HM, Aneja S. Perspectives of patients about artificial intelligence in health care. JAMA Netw Open. 2022; 5(5):e2210309. doi:10.1001/jamanetworkopen.2022.10309.

56. IntuitionLabs. AI adoption trends in US hospitals. 2026. Available from: https://intuitionlabs.ai/articles/ai-adoption-us-hospitals-trends.

57. US Food and Drug Administration. Evaluating real-world performance of AI-enabled devices. 2025. Available from: https://www.fda.gov/medical-devices/digital-health-center-excellence/request-public-comment-measuring-and-evaluating-artificial-intelligence-enabled-medical-device.

58. Topol EJ. As artificial intelligence goes multimodal, medical applications multiply. Science. 2023; 381(6663):adk6139. doi:10.1126/science.adk6139.

59. Aidoc. FDA breakthrough designation for CARE foundation model. 2025. Available from: https://www.aidoc.com/about/news/aidoc-receives-fda-breakthrough-device-designation.

60. Flower AI. Eye2Gene: federated learning in healthcare. 2025. Available from: https://flower.ai/blog/2025-03-07-eye2gene-revolutionizing-healthcare-with-federated-ai.

AI in Emergency Medicine: What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution, First Edition. Harvey Castro.
© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

5
AI in Surgery and Imaging: Precision on a Human Timeline

My mother was working three jobs when she saved her own life.

She had become an ultrasound technician, one of those jobs that immigrant mothers take because it pays decently and because she was too stubborn to stop climbing. One day, during a quiet moment between patients, she turned the probe on herself. She found a small nodule in her thyroid. Tiny. The kind of thing you could dismiss, the kind of thing a busy person working 70-hour weeks might decide to worry about later.

She didn't wait. She went to the doctor immediately. And the doctor found not one cancer, but two, one in her thyroid, one in her throat. Both deadly. Both treatable only because she caught them early. The margin between survival and the alternative was measured in weeks, not years. It was time she would not have had if she hadn't scanned herself that afternoon.

I think about this constantly. My mother survived because she happened to have access to an imaging machine and the training to use it. That is not a healthcare system. That is luck dressed up as competence. What about every other mother working three jobs who doesn't happen to be an ultrasound technician? Who gets their screening when insurance allows it, if insurance allows it, on a timeline dictated by guidelines written for populations, not people?

This is where AI in imaging becomes personal for me. Consider my patient Maria, a school teacher, a mother of two, who went for a routine mammography screening. The human reading cleared her. Normal. Go home. But that hospital had deployed an AI second opinion for mammography, and the algorithm flagged something the radiologist missed. They called Maria back. The biopsy confirmed early-stage breast cancer. She is fine today. She is fine because a machine caught what a human eye did not.

I cannot help but ask: if that AI second opinion had been available when my mother scanned herself, would it have caught the cancer even earlier? Would it have found the second cancer before the doctor did? I don’t know. That’s an honest answer. But I know this: early detection is not an abstract concept in my family. It is the reason my mother is alive. It is the reason she was sitting in the audience when I gave my TEDx talk, cancer-free, watching her son tell her story to a room full of strangers.

Every algorithm that improves imaging accuracy by even a fraction of a percent represents someone’s mother who might not have to save her own life by accident.

Executive Summary

There is a moment in every emergency department shift, and I have stood in it hundreds of times, when a patient arrives on a gurney, unable to speak, with an ischemic stroke evolving by the second. The window to intervene is narrow. The imaging suite is down the hall. The neurovascular specialist is across town or in another state. The margin between a full recovery and permanent disability is measured in minutes. In 2018, that margin began to shrink. Not because we added more physicians or faster ambulances, but because an algorithm learned to read a brain scan, recognize a blocked vessel, and fire a text message to the right surgeon in less than six minutes. What Viz.ai did for stroke that year is what this chapter is about: not the replacement of human judgment but its amplification, precisely when amplification matters most.

This chapter covers two rooms where AI is changing medicine at unprecedented speed: the operating room (OR) and the reading room. They share a common inflection point, roughly a decade in the making, when artificial intelligence crossed from research curiosity to clinical tool that patients are already benefiting from, and at which healthcare leaders must decide how to govern, deploy, and sustain it. As of December 2024, the FDA had authorized more than 1,016 AI-enabled medical devices, with radiology accounting for more than three-quarters of that total. The surgical robotics market, valued at approximately $12 billion in 2024, is projected to reach $27 billion or more by 2030. Intuitive Surgical alone placed 1,721

robotic systems in 2025 and crossed 20 million patients treated on its platforms. These are not pilot programs. They are infrastructure.

The most rigorous trial evidence, from the REAL trial for rectal cancer surgery and the MASAI trial for mammography screening, shows that AI-assisted systems outperform unaided practice in measurable clinical outcomes. The REAL trial demonstrated a 44 percent reduction in locoregional recurrence with robotic surgery compared to laparoscopy. The MASAI trial, the largest randomized controlled trial ever conducted in breast cancer screening, found a 29 percent increase in cancer detection and a 44 percent reduction in radiologist workload, with false-positive rates essentially identical to standard double reading.

The same data demands intellectual honesty. Across a full analysis of 86 deep-learning radiology algorithms, 81 percent degraded in accuracy when tested outside their training institution. A landmark study published in *The Lancet Gastroenterology & Hepatology* in 2025 documented the first real-world clinical evidence that AI tools can erode the very skills they are meant to support: colonoscopists who relied on AI assistance for three months saw their unassisted adenoma detection rate fall by 21 percent. In a sobering controlled experiment, radiologists given incorrect AI predictions dropped their diagnostic accuracy from 92.8 percent to 23.6 percent, overwhelmed by an algorithm they trusted too much.

This is the central paradox of AI in surgery and imaging: its greatest strength, consistent and tireless pattern recognition, is inseparable from its greatest risk, which is the normalization of deference. When the algorithm is right, it’s extraordinarily useful. When it’s wrong and humans have stopped checking, the consequences can be severe.

The title of this chapter is deliberate: “Precision on a Human Timeline.” Every surgical robot operating today is under surgeon control. Every cleared AI imaging tool functions as an adjunct, a second reader, a flagging system, a workflow accelerator, not an autonomous decision-maker. The question this chapter addresses is not “when will AI replace the surgeon or the radiologist?” because the answer, for any realistic planning horizon, is that it will not. The real question is: how do healthcare leaders, clinicians, and institutions capture the genuine benefits of AI while managing its

genuine risks, on a human timeline that keeps accountability where it belongs?

Throughout, I draw on my background as an emergency medicine physician and a student of both clinical data and AI systems. The patients in these stories are not abstractions (see Figure 5.1). They are the stroke victim whose outcome changed because a computer read her scan first. They are the veterans at VA medical facilities whose colonoscopies are now aided by an FDA-authorized polyp detection system. They are the rural patient who is significantly less likely to receive robotic surgery than his urban counterpart, a disparity this chapter names explicitly. Precision that serves only some patients is not yet the revolution it claims to be.

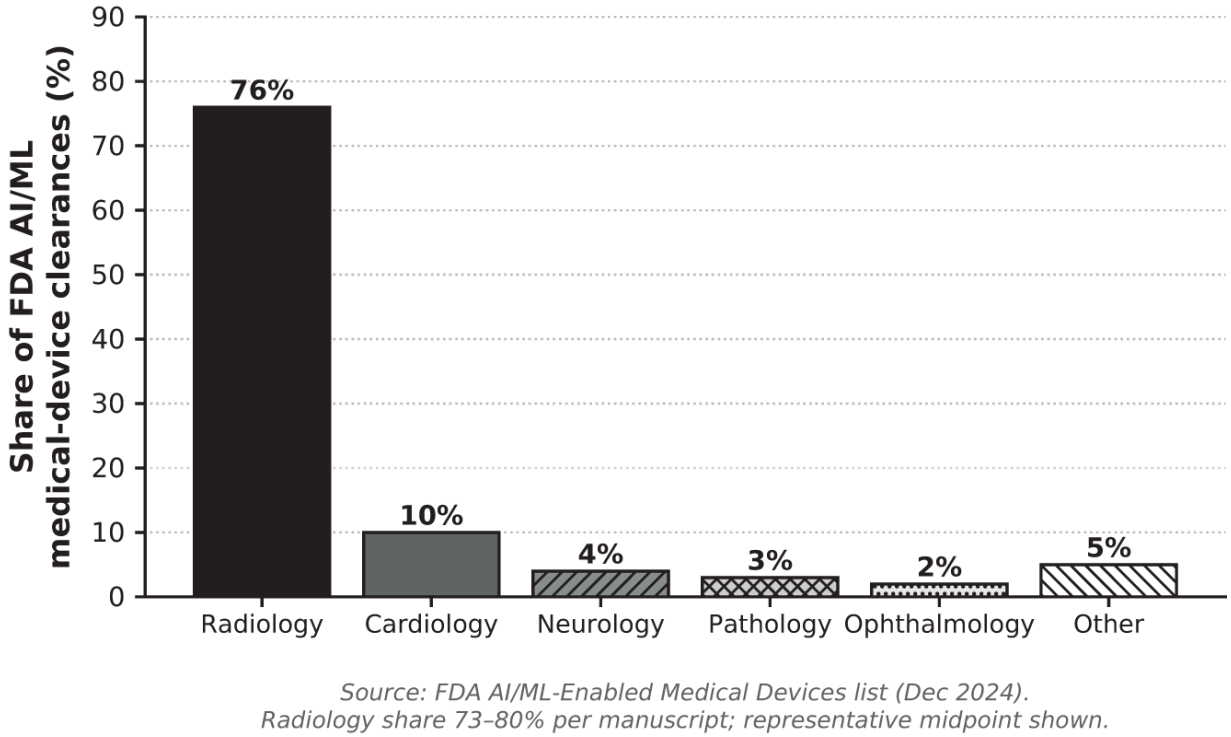


Figure 5.1 AI applications across six major medical specialties

Current Evidence

The evidence base for AI in specialized clinical practice has grown fast over the past five years. What follows covers the areas where peer-reviewed trials and real-world deployments have produced measurable clinical results.

The Operating Room: Surgical Robotics and AI-Assisted Systems

Robotic surgery in its current form is not autonomous. It's teleoperated: a surgeon's hand movements translated, filtered, and magnified by a mechanical system. The intelligence layered onto that mechanical substrate is growing, though, and the distinction between "robot that extends the surgeon's hands" and "system that assists the surgeon's decisions" is increasingly meaningful.

The Da Vinci Platform and the State of the Field

By the end of 2025, more than 12,000 Intuitive Surgical systems were in the global installed base, having collectively treated more than 20 million patients [1]. Annual procedure volume grew 18 percent year-over-year in 2025 to approximately 3.15 million procedures, and the company generated approximately $10.06 billion in revenue [1]. The da Vinci 5 (Model Is5000), cleared in March 2024, introduces capabilities that meaningfully advance AI integration. It contains more than 150 design innovations and 10,000 times the computing power of its predecessor [3]. Most clinically significant is its Force Feedback technology, the first of its kind among commercial surgical robots, where sensors at instrument tips detect push-and-pull forces and relay tactile sensation to the surgeon's hand controls. In preclinical testing, surgeons exerted up to 43 percent less force on tissue compared to the prior generation [3]. The platform also incorporates the My Intuitive+ analytics ecosystem capturing synchronized force, video, and motion data per case, giving surgeons AI-driven performance tracking through its Case Insights module [3]. By the end of 2024, 362 da Vinci 5 systems had been placed, with placement accelerating to 870 systems in 2025 alone [1,2].

The competitive field includes CMR Surgical's Versius Plus (510(k) clearance in December 2025, the world's second most-used soft-tissue robotic system with more than 40,000 procedures globally) [4], Medtronic's Hugo Ras (FDA clearance in December 2025 for urologic procedures covering approximately 230,000 annual U.S. procedures) [5], and Johnson & Johnson's Ottava (first U.S. clinical cases in April 2025 at Memorial Hermann-Texas Medical Center, De Novo authorization submission

expected early 2026) [6]. A 2024 systematic review found that 86 percent of FDA-cleared surgical robots operate at Level 1 (continuous surgeon control), 8 percent at Level 2 (pre-programmed task autonomy), and 6 percent at Level 3 (conditional autonomy). No cleared system yet achieves Level 4 or 5 full autonomy [7]. The global surgical robotics market was valued at approximately 11.98 billion dollars in 2024 and is projected to reach 27.14 billion dollars by 2030, at a compound annual growth rate of approximately 14.7 percent [8].

Clinical Outcomes Evidence: What Randomized Trials Show

The REAL trial, a multicenter RCT across 11 Chinese centers enrolling 1,240 patients comparing robotic versus laparoscopic total mesorectal excision for rectal cancer, stands as the strongest evidence of long-term oncological benefit. At three years, robotic surgery produced a locoregional recurrence rate of 1.6 percent versus 4.0 percent for laparoscopy (hazard ratio 0.45, $p = 0.03$), three-year disease-free survival of 87.2 percent versus 83.4 percent ($p = 0.04$), and perioperative complication rate of 16.1 percent versus 22.9 percent ($p = 0.003$). Functional outcomes, urinary, sexual, and defecatory, were significantly better in the robotic arm through the first postoperative year [9]. This is the first high-quality RCT evidence that robotic surgery may provide a long-term oncological advantage over laparoscopy in a specific indication.

The earlier ROLARR trial (471 patients, international centers) found no significant reduction in conversion rate (8.1 percent versus 12.2 percent, $p = 0.16$), longer operative time (+51 minutes, $p = 0.001$), and higher cost (+$1,132 per case, $p = 0.02$) with older-generation robotic systems. That divergence from REAL likely reflects technological generation differences [9]. A large retrospective cohort study of 19,769 patients undergoing colorectal surgery [10] found robotic surgery produced length of stay of 5.7 days versus 7.7 days laparoscopic and 11.3 days open, with mortality of 0.6 percent versus 1.0 percent and 4.2 percent, respectively, and complication rates of 3.7 percent versus 4.5 percent and 9.4 percent. Robotic versus laparoscopic mortality and complication differences were not significant after full risk adjustment, consistent with the broader finding that robotic

surgery's most unambiguous advantage is over open surgery, not laparoscopic [10].

A meta-analysis of 25 studies on AI-assisted robotic surgery in the *Journal of Robotic Surgery* reported directional aggregate findings of 25 percent reduction in operative time, 30 percent decrease in intraoperative complications, and 40 percent improvement in surgical precision, interpreted as directional rather than definitive given study heterogeneity [11]. In robotic cardiac surgery, pooled in-hospital mortality of 0.9 percent and overall graft patency of 96.4 percent for totally endoscopic coronary artery bypass grafting was comparable to conventional approaches in propensity-matched analyses [12].

Approaching Autonomy: STAR and SRT-H

The Smart Tissue Autonomous Robot (STAR) performed laparoscopic intestinal anastomosis on four pigs without continuous human guidance in January 2022, the first robot to perform laparoscopic surgery on soft tissue with minimal human intervention, producing fewer mistakes and more consistent suture spacing than expert surgeons [13]. In July 2025, Surgical Robot Transformer-Hierarchical (SRT-H) performed eight complete autonomous cholecystectomies on pig cadavers, all ex vivo, with 100 percent accuracy across 17 sequential surgical tasks, adapting in real time to unexpected events and responding to voice commands [14]. Both STAR and SRT-H are strictly research systems. No autonomous surgical robot has operated on a living human patient in any clinical setting as of March 2026, and researchers estimate more than a decade before such systems reach routine clinical practice [14].

The Reading Room: AI in Medical Imaging

The FDA had authorized 1,016 AI-enabled medical devices as of December 2024, growing to more than 1,356 by September 2025 [15]. Radiology commands approximately 73 to 80 percent of all medical AI authorizations. Between 94 and 97 percent are cleared via the 510(k) pathway [16,17].

Chest X-Ray AI Lunit Insight Cxr, trained on more than 1.2 million chest X-ray images, detects 11 major thoracic abnormalities with AUC of 0.932 in independent comparative testing [18]. AI assistance improved sensitivity

for all four major target findings, pneumonia, nodule, pleural effusion, and pneumothorax, without decreasing specificity. Pneumothorax sensitivity improved from 0.792 without AI to 0.965 with AI, while reporting time fell 10 percent [18]. Qure.ai's qXR, FDA-cleared and validated across 213,459 chest X-rays, demonstrated 96 to 99 percent specificity and 95 to 100 percent sensitivity for lung nodule detection, outperforming radiologist sensitivity of 74 to 76 percent for the same task, with a negative predictive value of 98.9 percent in a prospective multicenter study of 65,604 chest X-rays [19]. As of February 2026, Qure.ai holds 26 FDA-cleared indications across nine products, and its qXR-Detect product is the only chest X-ray CADe device cleared with a Predetermined Change Control Plan (PCCP), allowing algorithm updates without full re-clearance [19].

Mammography AI and the MASAI Trial

The MASAI trial, randomizing more than 105,000 women in Sweden with results published in *The Lancet* in January 2026, is the largest RCT ever conducted in cancer screening using AI. The AI arm (Transpara 1.7.0 plus one radiologist) detected 29 percent more cancers than standard double reading by two radiologists [20]. The primary endpoint (interval cancer rate) was 1.55 per 1,000 women versus 1.76 in the control arm, meeting noninferiority [20]. The AI arm produced 16 percent fewer invasive cancers, 21 percent fewer large cancers, and 27 percent fewer aggressive nonluminal-A cancers, with sensitivity of 80.5 percent versus 73.8 percent at identical specificity of 98.5 percent [20]. The AI arm also produced a 44 percent reduction in radiologist screen-reading workload at the 2023 interim analysis, with false-positive rates essentially identical (1.5 percent versus 1.4 percent) [20]. Lead researcher Dr. Kristina Lång noted: "Our study does not support replacing healthcare professionals with AI…but our results potentially justify using AI to ease the substantial pressure on radiologists' workloads" [20].

AI in Pathology

Paige Prostate Detect, the first FDA-authorized AI in digital pathology, cleared De Novo September 2021, improved pathologist sensitivity from 88.7 percent to 96.6 percent in its premarket study of 610 biopsy slides from 218 institutions. As a second reader, it reduced cancer detection errors

by 70 percent and false positives by 24 percent [21]. As of August 2025, only three AI/ML devices in digital pathology had received FDA clearance, reflecting a "dataset-driven validation gap" rather than absent regulatory pathways [22].

Ophthalmology AI IDx-DR (now LumineticsCore), the first autonomous AI diagnostic system cleared by the FDA (April 2018), demonstrated 87.2 percent sensitivity and 90.7 percent specificity in its registration trial across 900 subjects at ten primary care sites, with 100 percent negative predictive value in a 1,141-patient real-world validation [23]. EyeArt (cleared 2020) demonstrated 96 to 97 percent sensitivity across 942 subjects, with 91.3 percent real-world sensitivity in 101,710 consecutive patient visits [24]. CPT code 92229, established in 2021, enabled Medicare reimbursement for autonomous retinal analysis [24].

Radiologist-AI Collaboration

The strongest evidence consistently shows radiologist-AI collaboration outperforming either alone. A Moscow multicenter study of 3,409 brain CT studies found that radiologists using AI achieved 98.91 percent sensitivity and 99.83 percent specificity, versus AI alone at 95.91 percent sensitivity and 87.35 percent specificity, with a 73-fold reduction in false positives. All 12 cases missed by radiologists were caught by AI, and all 45 cases missed by AI were caught by radiologists [25].

Emerging Technologies: Foundation Models and Federated Learning

MedSAM, fine-tuned on more than 1.57 million medical image-mask pairs across 10 modalities and 30 cancer types, outperforms both the original SAM and specialist models across 86 internal and 60 external validation tasks [26].

BiomedParse (Microsoft Research, November 2024) unifies recognition, detection, and segmentation across nine modalities via natural-language prompt, outperforming MedSAM by 75 to 85 Dice score points on previously unseen evaluation data (held-out test sets) [27]. The Dice score measures how closely an AI's predicted boundary overlaps with the true boundary in medical image segmentation, where 100 represents a perfect

match. Scores of 75 to 85 on held-out test sets indicate strong but not yet expert-level performance. The German RACOON initiative showed multisite federated deployment for imaging AI, concluding that “despite substantial hurdles, the benefits clearly justify the investment” [28]. Adoption of AI in radiology varies dramatically: 48 percent of European radiologists use AI tools in routine practice, compared to an estimated 2 percent of U.S. radiology practices [29].

In my practice, the Moscow multicenter numbers resonate deeply: I’ve watched radiologists catch what the algorithm missed, and I’ve watched the algorithm catch what the radiologist missed. That complementarity is exactly why the “AI versus radiologist” framing is the wrong question.

Real-World Case Studies

These case studies show how AI tools have moved from controlled research settings into routine clinical workflows, with measurable effects on patient outcomes and operational efficiency.

Case 1: Viz.ai and the Race Against Stroke

On February 13, 2018, the FDA granted De Novo marketing clearance to Viz.ai’s ContaCT (Viz LVO), the first AI-based clinical decision support software cleared for sale in the United States for any condition [30]. Before Viz.ai, CT angiogram images entered a queue, a radiologist read them (sometimes minutes later, sometimes much later depending on staffing and workflow), and the stroke team mobilized when communication finally happened. After Viz.ai, the computed tomography angiography (CTA) images are analyzed by a convolutional neural network in real time. If a large vessel occlusion is detected, a text notification fires automatically to the neurovascular specialist while the first-line provider is still examining the patient. The median scan-to-notification time in FDA submission data? Less than six minutes. In 95 percent of cases, automated notification preceded standard-of-care communication by an average of 52 minutes [30].

The VALIDATE study, which surveyed more than 14,000 patients in 166 facilities in 17 states, found Viz LVO implementation reduced the median

door-to-neurointerventionalist notification time from 89.5 minutes to 50 minutes, a 39.5-minute statistically significant reduction ($p < 0.001$) [31]. A multicenter retrospective analysis at ISC 2025 found a 31-minute average reduction in treatment time across 474 patients [32]. A meta-analysis of 12 studies involving 15,595 patients confirmed that Viz.ai implementation significantly shortened every measured workflow interval [33]. One material caveat: the pooled analysis did not show statistically significant improvements in mortality or functional independence at 90 days, though the trend toward improved functional outcome (RR 1.31, $p = 0.08$) is encouraging [33].

As of January 2026, Viz.ai is deployed in nearly 2,000 hospitals, serving five patients per minute and backed by more than 120 peer-reviewed publications [34]. It was the first software to receive a Medicare New Technology Add-on Payment (NTAP)—an inpatient reimbursement for up to 65 percent of the technology's cost on top of standard DRG payments [34]. As stroke director James Siegler, MD, testified at ISC 2025: "Every one-minute delay to endovascular therapy has been associated with four additional days of disability-adjusted life-years" [32].

Case 2: GI Genius and the Colonoscopy Polyp Problem

Colorectal cancer is one of the most preventable cancers in medicine, yet adenoma miss rates in standard colonoscopy have historically ranged from 20 to 26 percent. On April 9, 2021, the FDA granted De Novo authorization to GI Genius, the first AI system cleared for colonoscopy in the United States [35]. The technology is simple in concept: a computer vision system processes live colonoscopy video in real time and marks potential polyps with a green square and a low-volume alert. The endoscopist sees everything they would normally see plus every location the AI flags, in real time, at no reduction in withdrawal speed.

The first U.S. trial, the DETECT study published in *Gastroenterology* in March 2022, enrolled 1,174 patients at eight centers and found adenoma miss rate fell from 32.4 percent with standard colonoscopy to 15.5 percent with GI Genius, a 50 percent reduction in missed colorectal polyps [35]. The COLO-DETECT trial, the largest pragmatic RCT of GI Genius, enrolling 2,032 participants at 12 NHS Trusts in England, found adenomas detected in 56.7 percent of AI-assisted patients versus 48.4 percent of

standard patients (adjusted OR 1.47, $p < 0.0001$), with comparable adverse events [36]. A meta-analysis of seven RCTs involving 9,639 patients confirmed a 12 percent relative increase in adenoma detection rate and a 27 percent relative increase in sessile serrated lesion detection. For approximately every 8 to 10 patients undergoing colonoscopy with GI Genius, one additional adenoma is detected that would otherwise have been missed [37].

In October 2022, Medtronic received a contract to install 115 GI Genius systems at VA medical facilities across the United States [38]. One real-world implementation caveat worth noting: a three-month institutional trial at one outpatient endoscopy unit found no statistically significant Adr improvement and discontinued the device. RCT-level benefits require trained endoscopists who reliably respond to device prompts [38].

Case 3: AI Catching What Humans Miss

Between August 2019 and April 2020, Sutter Health deployed an AI peer review platform across 33,185 CT scans, finding 210 missed lung nodules of which 123 were significant enough to change patient care, approximately eight times the detection rate of standard peer review [39]. Sutter Health estimated system-wide deployment would identify approximately 3,000 missed nodules and 120 malignancies over five years [39]. The Seoul National University Hospital RCT randomized 10,476 patients and found AI-assisted interpretation detected actionable lung nodules at 0.59 percent versus 0.25 percent without AI, more than double the detection rate, with no difference in false-referral rates [40].

A 2025 *Diagnostics* study of breast cancer cases found that lesions missed by AI were smaller and less calcified than those missed by human readers, confirming error complementarity: AI and human readers catch genuinely different cancers, making collaboration stronger than either alone [41].

Case 4: AI in Global Health (Kenya, Bangladesh, Ghana)

In Kenya, a pilot program used AI to analyze blood smear photos from a $50 portable microscope clipped to a smartphone, achieving 98.5 percent accuracy in diagnosing Plasmodium falciparum malaria in approximately 90 seconds. The pilot expanded to more than 420 health facilities across

eight counties, with early *The Lancet Digital Health* results showing a 31 percent reduction in inappropriate antibiotic prescribing and a 19 percent drop in severe malaria complications [42].

In Bangladesh, an independent evaluation published in *The Lancet Digital Health* compared five commercial AI algorithms for tuberculosis detection on 23,954 individuals at three screening centers. All five significantly outperformed radiologists. Two, Qure.ai's qXR and Delft Imaging's CAD4TB, met the WHO Target Product Profile threshold (90 percent sensitivity, 70 percent specificity), and all five reduced confirmatory Xpert tests required by 50 percent while maintaining greater than 90 percent sensitivity [43]. In March 2021, WHO updated its guidelines to recommend computer-aided detection software in place of human readers for digital chest X-ray analysis in TB screening [43].

In Ghana, Chestify AI, deployed across 25 health facilities, reduces diagnostic turnaround times by approximately 40 percent, with reports delivered within three hours instead of days [42]. Rwanda and Ghana are bundling AI diagnostics into national health insurance schemes at approximately one to two dollars per test [42]. Sub-Saharan Africa carries 24 percent of the global disease burden with 3 percent of the world's health workers. For these populations, AI is not a convenience: it is increasingly the only realistic path to timely diagnosis [42].

Risks and Limitations

The case for AI in surgery and imaging is real and growing. Intellectual honesty demands that the risks and limitations receive equal rigor.

Surgical Robot Adverse Events

The most thorough analysis of surgical robot adverse events, a 2016 PLOS One study using natural language processing on 14 years of FDA MAUDE data covering approximately 1.75 million robotic procedures, identified 144 deaths, 1,391 patient injuries, and 8,061 device malfunctions across 10,624 adverse event reports [44]. The rate of injury and death events remained approximately 83.4 per 100,000 procedures (95% CI: 74.2–92.7) since 2007, meaning the rate did not worsen as adoption grew, but absolute

numbers increased proportionally [44]. Device malfunction types included burnt or broken instrument pieces falling into patients (14.7 percent, causing 119 injuries and one death), electrical arcing causing tissue burns (10.5 percent, 193 injuries), and unintended instrument operation (8.6 percent, 52 injuries and two deaths) [44].

In 2025, a Reuters investigation reported that since AI was added to the TruDi Navigation System, the FDA received reports of at least 100 malfunctions and adverse events, compared to only seven prior to AI integration, with at least 10 patients sustaining injuries including strokes from inadvertent arterial damage [44].

Automation Bias: When AI Makes Humans Worse

Automation bias, the tendency to over-rely on algorithmic outputs and defer to them even when wrong, is the most underappreciated risk in AI-assisted medical imaging. In my experience, this is also one of the hardest things to train people out of. In a multisite prospective study of 220 physicians (Radiology, 2024), participants who received correct AI advice achieved diagnostic accuracy of 92.8 percent. When AI advice was incorrect, accuracy dropped to 23.6 percent [45]. The senior author stated plainly: “When we rely too much on whatever the computer tells us, that’s a problem, because AI is not always right” [45].

A study in *European Radiology* study using a sham AI system found that radiologist false negatives rose from 2.7 percent without AI to 20.7–33.0 percent when AI incorrectly indicated a normal finding. False positives rose from 51.4 percent to 80.5–86.0 percent under incorrect AI flagging. Radiologists were less likely to override AI errors when the AI result was kept in the patient’s record, suggesting that liability concerns around documented disagreement with an AI may reinforce the bias [46].

Generalizability Failures

A systematic review of 86 deep-learning radiology algorithms found that 81 percent showed decreased accuracy on external datasets, with nearly a quarter experiencing a drop of 0.10 AUC or more when tested outside their training institution [48]. A model trained on progressive interstitial pulmonary fibrosis cases dropped from approximately 100 percent cross-

validation accuracy to 45.3 percent sensitivity when tested on data with different patient positioning [48]. In AI mammography, a 2025 study in *Radiology* study evaluating commercial AI on 1,082 consecutive invasive breast cancers found the AI missed 14 percent overall, with a 17.2 percent false-negative rate for hormone-receptor-positive cancers, which represent approximately 70 percent of all breast cancers [47,49].

Deskilling: The Evidence We Cannot Ignore

A multicenter observational study provided the first real-world clinical evidence of AI-induced deskilling. Nineteen experienced colonoscopists at four Polish centers, each with more than 2,000 lifetime colonoscopies and an average of 27.6 years of experience, used AI-assisted colonoscopy for approximately three months. Their unassisted adenoma detection rate then fell from 28.4 percent to 22.4 percent, a six-percentage-point absolute decline and a 21 percent relative decline [50]. Only four of 19 endoscopists showed no decline. Researchers called this "the first documentation of a potential deskilling effect from clinical AI." The observational design and potential confounders, including a doubling of procedure volume over the period, require acknowledgment, but the pattern is consistent enough to warrant caution [50].

In robotic surgery, negative skill transfer from robotic to laparoscopic surgery has been documented in multiple studies. Primary robotic experience impaired subjects' laparoscopic task performance, while the reverse transfer, from laparoscopy to robotics, is positive [51]. Robotic surgical skills, meaning proficiency at the robotic console itself, degraded significantly within four weeks of inactivity in newly trained surgeons, with significant increases in task time and error rates at the four-week mark [51].

Cost Barriers and Access Inequity

The da Vinci 5 carries a purchase price of $1.8 to 2.5 million, approximately 30 percent above the predecessor Xi, with annual maintenance of 100,000 to 200,000 dollars and disposable per-procedure instrument costs averaging approximately 1,866 dollars per case [52]. A 2026 systematic review of 38 randomized trials found that conventional laparoscopy was more cost-favorable in the majority of head-to-head comparisons [53].

These costs create structural inequity. Rural patients were significantly less likely to undergo robotic lobectomy compared to urban patients [54]. Medicaid or uninsured patients were substantially less likely to receive robotic prostatectomy (OR 0.56, 95% CI: 0.45–0.70) compared to privately insured patients, even after controlling for hospital volume [55]. Because robotic surgery carries measurable outcome advantages over open surgery, access disparities translate directly into outcome disparities [55]. This is where the rubber meets the road: a technology that improves outcomes only for well-insured urban patients is not yet the improvement it's marketed as being [62].

Cybersecurity: A Real but Undercharacterized Risk

In 2015, University of Washington researchers became the first to successfully hack a teleoperated surgical robot, the Raven II, showing the ability to override commands, misuse the emergency stop, cause erratic movement, and take complete control of the device. The Interoperable Telesurgery Protocol was completely open and publicly available, and Raven II video feeds were broadcasting over the Internet without encryption [56]. The FBI's Cyber Division has reported that 53 percent of networked medical devices have at least one known critical vulnerability, 21 percent are secured by weak or default credentials, and 8 percent of imaging systems carry known exploited vulnerabilities linked to ransomware [57]. The 2023 State of Cybersecurity for Medical Devices report specifically flagged surgical navigation technologies as active threat targets [57].

As of the research date for this chapter, no publicly confirmed real-world cyberattack on a surgical robot has caused direct patient harm. The UW Raven II hack was a controlled research demonstration. Structural vulnerabilities are documented, though, and the consequences of a real-world mid-procedure system lockout or ransomware attack would not be theoretical.

Racial and Demographic Bias in AI Imaging

AI imaging tools can encode and replicate the biases present in their training data. An algorithm trained primarily on data from one demographic group may perform systematically worse on underrepresented populations,

an issue of both technical and ethical weight. Obermeyer et al. showed that a widely deployed commercial algorithm used to allocate healthcare resources had significant racial bias, assigning Black patients lower risk scores than equally sick white patients because it relied on healthcare spending as a proxy for health needs. Though this algorithm was not imaging-specific, the mechanism applies directly to medical imaging AI trained on non-representative datasets [58]. FDA guidance now explicitly recommends demographic subgroup performance reporting, but enforcement mechanisms remain limited and real-world performance monitoring across patient populations is not yet standard [66].

What I tell hospital executives is that these numbers don't lie: a tool that performs well in a predominantly white training dataset will fail quietly in a diverse patient population, and you won't know until the harm is done.

Regulatory Fragmentation and the Cleared-but-Not-Validated Problem

FDA clearance via the 510(k) pathway requires substantial equivalence to a predicate device, not independent proof of clinical benefit. In practice, a chest X-ray AI can be cleared based on technical equivalence to a prior algorithm without a prospective, multicenter clinical trial showing improvement in patient outcomes. The AI in Medical Imaging report in *Radiology* in 2022 found that fewer than 6 percent of cleared AI imaging algorithms had published evidence of clinical utility in peer-reviewed literature at the time of clearance [59]. This isn't a regulatory failure per se: the FDA's role is safety and substantial equivalence, not efficacy demonstration. It does mean that institutions deploying 510(k)-cleared imaging AI can't assume clinical benefit has been independently validated. Prospective local monitoring is not optional. It's a patient safety obligation.

Ethics, Equity, and Patient Rights

The rapid integration of AI into clinical workflows raises urgent questions about patient autonomy, transparency in algorithmic decision-making, and equitable access across diverse populations. Current regulatory frameworks have only begun to address these concerns.

Informed Consent and the Disclosure Gap

The ethical obligation to disclose AI involvement in clinical care is the subject of active and unresolved debate. The American Medical Association's policy on augmented intelligence holds that physicians must disclose when AI tools are involved in clinical decision-making that materially affects patient care. The American College of Radiology's Arch-AI program makes disclosure policy a governance requirement [60]. Practice is inconsistent. One multi-institution survey found that fewer than 20 percent of radiology practices had a formal policy requiring AI disclosure to patients, and most patients undergoing AI-assisted diagnostic imaging were unaware that AI was involved [60].

The disclosure gap is real, and the field hasn't resolved how detailed that disclosure needs to be. Should a patient know that an algorithm flagged the finding their radiologist confirmed? Probably. Should consent forms enumerate every AI tool in the workflow? That may not be practical or even useful to patients. The right approach is not blanket disclosure but thoughtful transparency focused where patients have meaningful choices to make.

Surgical Video Privacy

Surgical video data generated by robotic systems represents a category of medical information with significant privacy implications and limited regulatory clarity. Under HIPAA, surgical videos constitute protected health information (PHI). De-identified data can be used for AI training without patient authorization, but the "mosaic effect," the ability of modern AI to cross-reference de-identified records with public datasets to re-identify individuals, challenges the adequacy of Safe Harbor standards [67]. The emerging institutional governance consensus is that patients should be informed about video recording with an opportunity to consent or decline, access should be limited to predetermined use cases, and IRB oversight is appropriate for systematic video collection programs [67].

Patient Autonomy: The Right to Refuse AI

In the United States, no specific statute provides patients a right to refuse AI-assisted diagnosis or treatment. In the European Union, GDPR Article

22's right against solely automated decision-making is being analyzed as a potential foundation for AI refusal rights in healthcare, though member state health laws do not yet directly equip patients with legal means to refuse AI-aided decisions [65]. Scholars note a countervailing tension: excessive disclosure about AI use "can even harm patients, stoking distrust of clinicians' recommendations." The prescription is not less transparency but more thoughtful transparency, focused where patients have meaningful choices to make [65].

Leadership Implementation Framework

For CEOs, CIOs, and CMOs standing at the intersection of clinical opportunity and institutional risk, AI in surgery and imaging is neither a straightforward investment nor an avoidable one. The following framework reflects current best practice across regulatory, clinical, and operational dimensions.

Governance: ACR Arch-AI and Assess-AI

The ACR launched its ACR Recognized Center for Healthcare-AI (Arch-AI) program in June 2024, the first national AI quality assurance program for radiology facilities in the United States [68]. To achieve recognition, facilities must form an interdisciplinary AI governance group, maintain a documented inventory of deployed algorithms, comply with security protocols, review and document AI use cases and training procedures, monitor algorithm performance, and contribute data to the ACR Assess-AI central registry [68]. Assess-AI, launched November 2024, is the world's first AI quality registry designed to monitor ongoing real-world clinical performance of imaging AI algorithms, tracking concordance and discordance between AI results and radiologist reports and monitoring stability of algorithm inputs over time [68].

"Algorithmovigilance," systematic post-deployment monitoring of AI performance in real-world conditions, is not yet standard practice at most institutions. The 2025 FDA lifecycle guidance explicitly requires post-market drift monitoring, and the EU AI Act mandates performance monitoring plans. Building this capability before it is legally required is the responsible approach [64,68].

Credentialing: SAGES-MIRA and the FRS Gap

The Society of American Gastrointestinal and Endoscopic Surgeons (SAGES) and the Minimally Invasive Robotic Association (MIRA) have issued consensus credentialing guidelines requiring documentation of prior robotic experience, hands-on course completion, case observation and proctored cases, and training for the entire operative team, explicitly rejecting case-count-based credentialing in favor of proficiency-based assessment [69]. The critical gap: the American Board of Surgery requires Fundamentals of Laparoscopic Surgery (FLS) for residency graduation but does not require Fundamentals of Robotic Surgery (FRS) [69]. Hospital credentialing policies requiring an average of only 3.24 proctored cases, while the literature shows 19 to 128 cases to reach performance plateau in colorectal robotic surgery, represent a patient safety gap that institutional leaders should address in medical staff bylaws without waiting for national mandates [59].

Post-Deployment Monitoring

For surgical robotics, institutions should model capital acquisition cost (1.8 to 2.5 million dollars for da Vinci 5), annual service contracts (100,000 to 200,000 dollars), and per-procedure instrument costs averaging $1,866 against: reduced length of stay, reduced readmission and complication rates, and the volume threshold below which amortized costs become prohibitive (approximately 200 to 300 cases per system per year) [52]. A 2023 systematic review of 33 economic evaluations found robotic-assisted surgery was cost-effective or potentially cost-effective in 81 percent of studies, particularly for radical prostatectomy, nephrectomy, and cystectomy, when societal perspective, longer time horizons, and adequate case volumes were modeled [53].

For imaging AI, ROI modeling should include radiologist time savings (documented at 36 percent workload reduction for chest X-ray AI and 44 percent for mammography), reduction in missed findings and associated downstream costs, and reimbursement through NTAP, TCET, and Category I CPT codes.

Training Requirements

Institutions should implement three-tier training requirements for AI-assisted surgery and imaging. At the individual clinician level: procedural training meeting or exceeding SAGES-Mira guidelines for surgery; AI tool-specific training with documented simulation performance for imaging; and periodic competency reassessment including unassisted performance testing to monitor for deskilling. At the team level: standardized OR team training for robotic procedures, including equipment failure protocols and immediate conversion procedures, because a surgeon who has grown dependent on robotic assistance must be able to manage the case when the robot fails. At the institutional level: an interdisciplinary AI governance committee with authority to approve, monitor, and discontinue deployed AI tools; a documented vendor evaluation process including independent generalizability assessment; and explicit policies for managing automation bias, including structured reading protocols requiring independent radiologist assessment before AI output review for high-stakes diagnoses.

Future Outlook

The trajectory in surgical robotics and clinical imaging AI points toward increasingly autonomous systems, though the path remains technically complex and ethically contested. Real-world deployment timelines are measured in years and decades, not months.

Autonomous Surgical Systems: The Real Timeline

The Johns Hopkins SRT-H system performing autonomous cholecystectomies on pig cadavers with 100 percent success in July 2025 represents genuine scientific progress, trained on only 34 demonstration videos and showing that surgical autonomy is technically achievable for defined, repeatable procedures [14]. Technically achievable in a research setting and clinically deployable in a human patient are separated by multiple regulatory, ethical, and safety milestones that will take well over a decade to clear.

No FDA regulatory framework for autonomous surgical systems yet exists. No Level 4 or Level 5 autonomous surgical robot has been cleared. The

clinical trial infrastructure required, across diverse patient anatomies, complication profiles, and operative scenarios, would demand years of animal studies, cadaveric studies, and carefully constructed first-in-human trials. Researchers consistently estimate the preclinical-to-clinical journey at more than a decade. Leaders who plan technology strategy around autonomous surgery as a near-term reality are planning for the wrong horizon.

Foundation Models for Imaging

The shift from task-specific AI algorithms, trained to detect one finding in one modality, to foundation models trained on massive multimodal datasets represents the next structural change in imaging AI. MedSAM, BiomedParse, and BiomedGPT show a model architecture that can perform across imaging modalities using natural-language prompts, without requiring task-specific fine-tuning [26,27]. The practical upside is an AI that a radiologist can query, “show me all focal liver lesions in this CT,” rather than a series of separate FDA-cleared tools each handling one finding. The FDA has not yet cleared any foundation model or large language model-based medical device as of September 2025, and the regulatory pathway for continuously-updating systems will require frameworks beyond current PCCP guidance [15].

Multimodal Integration and the Converging Room

The surgical operating room and the radiology reading room are converging. Proprio’s FDA-cleared surgical guidance platform integrates preoperative imaging into real-time intraoperative visualization with continuous 3D anatomy overlays during spine surgery [70]. SimBioSys’s TumorSight Viz converts breast MRI into 3D surgical planning models with same-day results [70]. The da Vinci 5’s roadmap includes integration of preoperative imaging directly into the surgical field, anatomical structure highlighting, and tissue and cancer identification during the procedure [3].

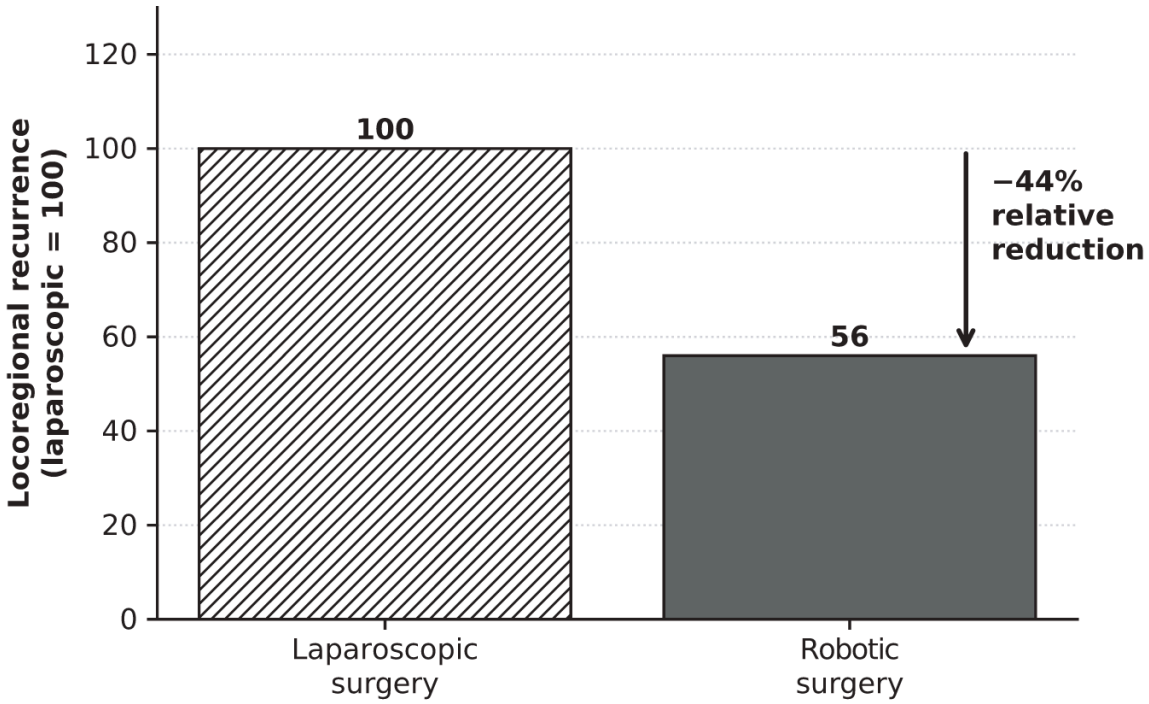
I’ve found this convergence trend to be one of the most clinically meaningful developments to follow: when the reading room’s insight and the operating room’s action are unified in real time, the surgeon operates with a level of contextual intelligence that simply didn’t exist before. The longer-term vision, multimodal AI integrating imaging data, EHR data,

genomics, intraoperative video, and physiological monitoring into a unified surgical intelligence layer, is technically plausible and commercially nascent. The trajectory is clear.

AI Surgical Coaching

One of the most promising near-term applications is AI as a surgical coach rather than a co-pilot. The Northwell Health and Lenox Hill Hospital team is developing AI guidance systems, estimated two to three years from commercial availability as of late 2025, that alert surgeons in real time when too close to critical structures or when a procedure diverges from expected workflow [71]. The Meta Segment Anything Model 2 has reduced manual annotation time for surgical video by 85 to 90 percent, dramatically accelerating development of training and feedback tools [71]. Surgical phase recognition systems already achieve 85 to 88 percent accuracy in recognizing procedure phases in real-time colonoscopy and hernia repair videos, forming the foundation for automated quality assessment, trainee feedback, and OR workflow optimization [71].

> **AI in surgery and imaging is already clinical, not experimental.** More than 1,016 AI-enabled medical devices had received FDA authorization as of December 2024, with radiology commanding 73 to 80 percent of all clearances. Intuitive Surgical’s da Vinci platform has treated more than 20 million patients and generated $10.06 billion in revenue in 2025 [1,15].



Source: REAL trial (rectal cancer), manuscript Chapter 5.

Figure 5.2 The oncology AI continuum: screening through survivorship

The strongest clinical evidence for robotic surgery shows genuine oncological advantage in specific indications. The REAL trial showed a 44 percent relative reduction in locoregional recurrence with robotic versus laparoscopic total mesorectal excision for rectal cancer (1.6 percent versus 4.0 percent, HR 0.45, p = 0.03). For the right procedures in the right hands, the evidence is not just comparable to laparoscopy. It's better [9].

AI in mammography screening increases cancer detection by 29 percent without increasing false positives, while reducing radiologist workload by 44 percent. The MASAI trial, 105,000 women and the largest RCT in cancer screening, showed an interval cancer rate non-inferior to standard double reading, with results in *The Lancet* in January 2026 [20] (see Figure 5.2).

Key Takeaways

Automation bias is a quantifiable patient safety risk. When radiologists received incorrect AI predictions in a 2024 *Radiology* study of 220 physicians, diagnostic accuracy dropped from 92.8 percent to 23.6 percent. Institutional protocols must require independent radiologist assessment before AI output review on high-stakes diagnoses [45].

Eighty-one percent of deep-learning radiology algorithms degrade in external validation. Every AI imaging tool deployed should be prospectively monitored in the local clinical environment. The ACR's Assess-AI registry is the infrastructure for doing this systematically [48].

The first real-world clinical evidence of AI-induced deskilling has been published. Experienced colonoscopists' unassisted adenoma detection rate fell 21 percent after three months of AI-assisted colonoscopy (28.4 percent to 22.4 percent). Training programs and credentialing processes should explicitly account for maintaining unassisted competency alongside AI-assisted practice [50].

Access to robotic surgery is deeply inequitable. Rural patients face an adjusted OR of 0.28 for robotic versus open lung surgery, and Medicaid patients face an OR of 0.56 for robotic versus open prostatectomy compared to privately insured patients. Because these modality differences carry clinical outcome differences, access disparities are outcome disparities [54,55].

No autonomous surgical system has operated on a living human patient in any clinical setting. The Johns Hopkins SRT-H showed autonomous cholecystectomy on ex vivo pig cadavers with 100 percent success in July 2025, but expert consensus places autonomous surgical systems in routine clinical practice more than a decade away. Technology strategy built on autonomous surgery as a near-term reality is built on the wrong foundation [14].

Reimbursement is the primary adoption gatekeeper for imaging AI. Only three AI solutions hold Category I CPT codes with established payment rates in the United States as of January 2026.

Twenty-three others exist only as Category III temporary codes that do not guarantee payment. Healthcare leaders who engage CMS policy processes now are shaping the trajectory of the field [61].

References

1. Intuitive Surgical. Intuitive announces preliminary fourth quarter and full-year 2025 financial results [Internet]. 2026. Available from: https://finance.yahoo.com/news/intuitive-announces-preliminary-fourth-quarter-140500533.html.

2. GlobalData. Intuitive Surgical expands platforms despite market headwind. Medical Device Network [Internet]. 2026 Jan 15 [cited 2026Jul 5]. Available from: https://www.medicaldevice-network.com/analyst-comment/intuitive-surgical-expands-platforms-despite-market-headwind.

3. Intuitive Surgical. Da Vinci 5 force feedback technology [Internet]. 2024. Available from: https://www.intuitive.com/en-us/about-us/newsroom/force-feedback.

4. CMR Surgical. CMR Surgical secures 510(k) clearance of its Versius Plus robotic surgical system [Internet]. 2025. Available from: https://us.cmrsurgical.com/news/cmr-surgical-secures-510k-clearance-of-its-versius-plus-robotic-surgical-system.

5. Kelly S. Medtronic's Hugo surgical robot earns FDA clearance. MedTech Dive [Internet]. 2025 Dec 4 [cited 2026 Jul 5]. Available from: https://www.medtechdive.com/news/Medronic-Hugo-surgical-robot-FDA-clearance/807049.

6. Johnson & Johnson. Johnson & Johnson MedTech announces completion of first cases with Ottava robotic surgical system [Internet]. 2025. Available from: https://www.jnj.com/media-center/press-releases/johnson-johnson-medtech-announces-completion-of-first-cases-with-ottava-robotic-surgical-system.

7. Lee A, Baker TS, Bederson JB, Rapoport BI. Levels of autonomy in FDA-cleared surgical robots: a systematic review. NPJ Digit Med. 2024;

7:103. doi:10.1038/s41746-024-01102-y.

8. MarketsandMarkets. Surgical robots market [Internet]. 2025. Available from: https://www.marketsandmarkets.com/Market-Reports/surgical-robots-market-256618532.html.

9. Geropoulos G, Massias S, Pajaziti Q, Thapa B, Mohsin SN, Patel V. What do randomised trials reveal about robotic surgery? J Clin Med [Internet]. 2025. Available from: https://pmc.ncbi.nlm.nih.gov/articles/Pmc12524502.

10. Vazquez VK, Mena Albors L, Liu H, Ang D. Outcomes comparing 19,769 patients who underwent robotic, laparoscopic, or open colorectal surgery. JSLS. 2025; 29:e2025.00038. doi:10.4293/JSLS.2025.00038. PMID 41311504.

11. Ng Kok Wah J. The rise of robotics and AI-assisted surgery in modern healthcare. J Robot Surg. 2025; 19:298. doi:10.1007/s11701-025-02485-0. PMID 40540146.

12. Hwang B, Ren J, Wang K, Williams ML, Yan TD. Systematic review and meta-analysis of two decades of reported outcomes for robotic coronary artery bypass grafting. Ann Cardiothorac Surg. 2024; 13:82-98. doi:10.21037/acs-2023-rcabg-0191.

13. Johns Hopkins Hub. Robot performs first laparoscopic surgery without human help [Internet]. Baltimore (MD): Johns Hopkins University; 2022 Jan 26 [cited 2026 Jul 5]. Available from: https://hub.jhu.edu/2022/01/26/star-robot-performs-intestinal-surgery.

14. Hale C. Johns Hopkins' AI-powered robot performs autonomous surgery. Fierce Biotech [Internet]. 2025 Jul 10 [cited 2026 Jul 5]. Available from: https://www.fiercebiotech.com/medtech/johns-hopkins-ai-powered-voice-controlled-robot-performs-autonomous-surgery.

15. Casey B. FDA AI approvals surge past 1k for radiology. The imaging wire [Internet]. 2025 Dec 10 [cited 2026 Jul 5]. Available from:

https://theimagingwire.com/2025/12/10/ai-enabled-medical-devices-granted-fda-marketing-authorization.

16. Meskó B. Current state: over 1,250 FDA-approved AI-based medical devices [Internet]. 2025. Available from: https://www.linkedin.com/pulse/current-state-over-1250-fda-approved-ai-based-medical-mesk%C3%B3-md-phd-itl6f.

17. Sivakumar R, Lue B, Kundu S. FDA approval of artificial intelligence and machine learning devices in radiology: a systematic review. JAMA Netw Open. 2025; 8:e2542338. doi:10.1001/jamanetworkopen.2025.42338.

18. Ahn JS, Ebrahimian S, McDermott S, Lee S, Naccarato L, Di Capua JF, et al. Association of artificial intelligence-aided chest radiograph interpretation with reader rerformance and Efficiency. JAMA Netw Open. 2022; 5:e2229289. doi:10.1001/jamanetworkopen.2022.29289.

19. Zavaleta-Monestel E, Quesada-Villaseñor R, Arguedas-Chacón S, García-Montero J, Barrantes-López M, Salas-Segura J, et al. Revolutionizing healthcare: qure. AI's innovations in medical diagnosis and treatment. Cureus. 2024; 16:e61585. doi:10.7759/cureus.61585. PMID 38962585.

20. Ecancermedicalscience. AI-supported mammography screening results in fewer aggressive and advanced breast cancers [Internet]. 2026. Available from: https://ecancer.org/en/news/27721-ai-supported-mammography-screening-results-in-fewer-aggressive-and-advanced-breast-cancers-finds-full-results-from-first-randomised-controlled-trial.

21. Paige. Paige publishes data supporting first FDA-authorized AI software for digital pathology [Internet]. 2023. Available from: https://www.paige.ai/press-releases/paige-publishes-data-supporting-first-fda-authorized-ai-software-for-digital-pathology.

22. Makhlouf HR, Ossandon MR, Farahani K, Lubensky I, Harris LN. Digital pathology imaging artificial intelligence in cancer research and

clinical trials: an NCI workshop report. J Pathol Inform. 2025; 20:100531. doi:10.1016/j.jpi.2025.100531. PMID 41476571.

23. Abràmoff MD, Lavin PT, Birch M, Shah N, Folk JC. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. npj Digit Med. 2018;1:39. doi:10.1038/s41746-018-0040-6. Available from: https://www.nature.com/articles/s41746-018-0040-6; Riotto E, Gasser S, Potic J, Sherif M, Stappler T, Schlingemann R, et al. Accuracy of autonomous artificial intelligence-based diabetic retinopathy screening in real-life clinical practice. J Clin Med. 2024;13(16):4776. doi:10.3390/jcm13164776. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC11355215.

24. Ipp E, Liljenquist D, Bode B, Shah VN, Silverstein S, Regillo CD, et al. Pivotal evaluation of an artificial intelligence system for autonomous detection of referable and vision-threatening diabetic retinopathy. JAMA Netw Open. 2021; 4:e2134254. doi:10.1001/jamanetworkopen.2021.34254.

25. Khoruzhaya AN, Sakharova PA, Arzamasov KM, Kremneva EI, Burenchev DV, Erizhokov RA, et al. Standalone AI versus AI-assisted radiologists in emergency ICH detection: a prospective, multicenter diagnostic accuracy study. J Clin Med. 2025; 14:5700. doi:10.3390/jcm14165700. PMID 40869525.

26. Ma J, He Y, Li F, Han L, You C, Wang B. Segment anything in medical images. Nat Commun. 2024; 15:654. doi:10.1038/s41467-024-44824-z.

27. Potts B. Introducing BiomedParse, a groundbreaking foundation model for biomedical image analysis. Microsoft Research Blog [Internet]. 2025 [cited 2026 Jul 5]. Available from: https://www.microsoft.com/en-us/research/blog/biomedparse-a-foundation-model-for-smarter-all-in-one-biomedical-image-analysis.

28. Bujotzek MR, Akünal Ü, Denner S, Neher P, Zenk M, Frodl E, et al. Real-world federated learning in radiology: hurdles to overcome and benefits to gain. J Am Med Inform Assoc. 2025; 32:193-205. doi:10.1093/jamia/ocae259.

29. IntuitionLabs. AI radiology trends 2025 review [Internet]. 2026. Available from: https://intuitionlabs.ai/articles/ai-radiology-trends-2025.

30. Imaging Technol News. FDA clears first AI-powered clinical decision support software for stroke [Internet]. 2018. Available from: https://www.itnonline.com/content/fda-clears-first-ai-powered-clinical-decision-support-software-stroke.

31. Viz.ai. Large real-world, multi-center study demonstrates Viz.ai Platform saves critical minutes in stroke care [Internet]. 2023 Feb 8 [cited 2026 Jul 5]. Available from: https://www.viz.ai/news/large-real-world-multi-center-study-demonstrates-viz-ai-platform-saves-critical-minutes-in-stroke-care.

32. Viz.ai. Two new studies demonstrate proven impact of Viz.ai's stroke solution on patient outcomes and hospital economics [Internet]. 2025 Feb 13 [cited 2026 Jul 5]. Available from: https://www.viz.ai/news/new-studies-demonstrate-impact-of-vizais-stroke-solution.

33. Sarhan K, Azzam AY, Moawad MHE, Serag I, Abbas A, Sarhan AE. Automated emergent large vessel occlusion detection using Viz.ai software and its impact on stroke workflow metrics and patient outcomes in stroke centers: a systematic review and meta-analysis. Transl Stroke Res. 2025; 16:1487-1503. doi:10.1007/s12975-025-01354-0. PMID 40335883.

34. Viz.ai. Viz.ai Closes 2025 with record scale and patient impact, achieving profitability in its healthcare business while accelerating life sciences growth [Internet]. 2026 Jan 12 [cited 2026 Jul 5]. Available from: https://www.viz.ai/news/viz-ai-closes-2025-with-record-scale-and-patient-impact.

35. Medtronic. U.S. FDA grants De Novo clearance for GI Genius intelligent endoscopy module [Internet]. 2021. Available from: https://news.medtronic.com/2021-04-12-U-S-FDA-Grants-De-Novo-Clearance-for-First-and-Only-Artificial-Intelligence-System-

for-Colonoscopy-Medtronic-Launches-GI-Genius-TM-Intelligent-Endoscopy-Module.

36. Seager A, Sharp L, Neilson LJ, Brand A, Hampton JS, Lee TJW, et al. Polyp detection with colonoscopy assisted by the GI Genius artificial intelligence endoscopy module compared with standard colonoscopy in routine colonoscopy practice (COLO-DETECT): a multicentre, open-label, parallel-arm, pragmatic randomised controlled trial. Lancet Gastroenterol Hepatol. 2024; 9(10):911-923. doi:10.1016/S2468-1253(24)00161-4. PMID 39153491.

37. Sattar A, Sattar A, Khan MH, Zahid M, Tariq S, Choudhary N, et al. Effectiveness of the GI genius computer-aided detection system versus standard colonoscopy: a systematic review and meta-analysis of randomized controlled trials. Cureus. 2025; 17:e94624. doi:10.7759/cureus.94624. PMID 41246784.

38. Medical Device Network. GI Genius deployment [Internet]. 2023. Available from: https://www.medicaldevice-network.com/projects/gi-genius-intelligent-endoscopy-module-us.

39. AuntMinnie. AI peer review can spot missed lung cancer [Internet]. 2020. Available from: https://www.auntminnie.com/imaging-informatics/artificial-intelligence/article/15626767/c-mimi-ai-peer-review-can-spot-missed-lung-cancer.

40. RSNA. AI Improves Lung Nodule Detection on Chest X-rays [press release; Internet]. Radiological Society of North America; 2023 Feb 7 [cited 2026 Jul 5]. Available from: https://www.rsna.org/media/press/i/2415.

41. van Winkel SL, Sechopoulos I, Rodríguez-Ruiz A, Veldkamp WJH, Gennaro G, Chevalier M, et al. An exploration of discrepant recalls between AI and human readers of malignant lesions in digital mammography screening. Diagnostics (Basel). 2025; 15:1566. doi:10.3390/diagnostics15121566. PMID 40564886.

42. Okolue CA. Researchers in Africa are using AI to fill the global health care gap. Global Voices [Internet]. 2026 Feb 5 [cited 2026 Jul 5].

Available from: https://globalvoices.org/2026/02/05/researchers-in-africa-are-using-ai-to-fill-the-global-health-care-gap.

43. Qin ZZ, Ahmed S, Sarker MS, Paul K, Adel ASS, Naheyan T, et al. Tuberculosis detection from chest x-rays for triaging in a high tuberculosis-burden setting: an evaluation of five artificial intelligence algorithms. Lancet Digit Health. 2021; 3(9):e543-e554. doi:10.1016/S2589-7500(21)00116-3.

44. Alemzadeh H, Raman J, Leveson N, Kalbarczyk Z, Iyer RK. Adverse events in robotic surgery: a retrospective study of 14 years of FDA data. PLoS One. 2016; 11(4):e0151470. doi:10.1371/journal.pone.0151470.

45. RSNA. Incorrect AI advice influences diagnostic decisions [Internet]. Radiological Society of North America News; 2024 Nov 19 [cited 2026 Jul 5]. Available from: https://www.rsna.org/news/2024/november/ai-influences-diagnostic-decisions.

46. Bernstein MH, Atalay MK, Dibble EH, Maxwell AWP, Karam AR, Agarwal S, et al. Can incorrect artificial intelligence (AI) results impact radiologists, and if so, what can we do about it? A multi-reader pilot study of lung cancer detection with chest radiography. Eur Radiol. 2023; 33:8263-8269. doi:10.1007/s00330-023-09747-1. PMID 37266657.

47. Yu F, Moehring A, Banerjee O, Salz T, Agarwal N, Rajpurkar P. Heterogeneity and predictors of the effects of AI assistance on radiologists. Nat Med. 2024; 30:837-849. doi:10.1038/s41591-024-02850-w.

48. Yu AC, Mohajer B, Eng J. External validation of deep learning algorithms for radiologic diagnosis: a systematic review. Radiol Artif Intell. 2022; 4(3):e210064. doi:10.1148/ryai.210064. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC9152694.

49. McKee J. How well does AI detect invasive breast cancers? [Internet]. Radiological Society of North America News; 2025 Oct 29 [cited 2026 Jul 5]. Available from:

https://www.rsna.org/news/2025/october/ai-false-negatives-in-breast-cancer.

50. Budzyń K, Romańczyk M, Kitala D, Kołodziej P, Bugajski M, Adami HO, et al. Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study. Lancet Gastroenterol Hepatol. 2025; 10(10):896-903. doi:10.1016/S2468-1253(25)00133-5. PMID 40816301.

51. Jenison EL, Gil KM, Lendvay TS, Guy MS. Robotic surgical skills: acquisition, maintenance, and degradation. JSLS. 2012; 16:218-228. doi:10.4293/108680812X13427982376185. PMID 23477169.

52. Kelly S. Intuitive details launch plan for new da Vinci 5 robot. MedTech Dive [Internet]. 2024 Mar 18 [cited 2026 Jul 5]. Available from: https://www.medtechdive.com/news/intuitive-da-vinci-5-robot-launch-plan/710547.

53. Sadri H, Fung-Kee-Fung M, Shayegan B, Garneau PY, Pezeshki P. A systematic review of full economic evaluations of robotic-assisted surgery in thoracic and abdominopelvic procedures. J Robot Surg. 2023; 17(6):2671-2685. doi:10.1007/s11701-023-01731-7. PMID 37843673.

54. Erhunmwunsee L, Bhandari P, Sosa E, Sur M, Ituarte PHG, Lui NS. Socioeconomic, rural, and insurance-based inequities in robotic lung cancer resections. Video-assist Thorac Surg. 2020; 5:26.

55. Logan CD, Mahenthiran AK, Siddiqui MR, French DD, Hudnall MT, Patel HD, et al. Disparities in access to robotic technology and perioperative outcomes among patients treated with radical prostatectomy. J Surg Oncol. 2023; 128:375-384. doi:10.1002/jso.27274. PMID 37036165.

56. Univ Wash. Security risks in teleoperated surgical robots [Internet]. 2015. Available from: https://www.ee.washington.edu/spotlight/the-first-to-hack-a-teleoperated-surgical-robot-uw-ee-researchers-prove-security-risks-exist.

57. C2A Security. Healthcare and medical device cybersecurity statistics [Internet]. 2025. Available from: https://c2a-sec.com/60-healthcare-and-medical-device-cybersecurity-risk-statistics-for-2025.

58. Gordon WJ, Ikoma N, Lyu H, Jackson GP, Landman A. Protecting procedural care—cybersecurity considerations for robotic surgery. NPJ Digit Med. 2022; 5:148. doi:10.1038/s41746-022-00693-8. PMID 36127420.

59. Tamhankar A, Spencer N, Hampson A, Noel J, El-Taji O, Arianayagam R, et al. Real-time assessment of learning curve for robot-assisted laparoscopic prostatectomy. Ann R Coll Surg Engl. 2020; 102:717-725. doi:10.1308/rcsann.2020.0139. PMID 32538121.

60. King & Spalding. FDA predetermined change control plan guidance for AI-enabled devices [Internet]. 2024. Available from: https://www.kslaw.com/news-and-insights/fda-publishes-final-predetermined-change-control-plan-guidance-for-ai-enabled-device-software-functions.

61. Bipartisan Policy Center. Paying for AI in U.S. health care [Internet]. 2026. Available from: https://bipartisanpolicy.org/issue-brief/paying-for-ai-in-u-s-health-care.

62. Medicare.org. Does Medicare cover robotic surgery? [Internet]. Available from: https://www.medicare.org/articles/does-medicare-cover-robotic-surgery.

63. Penn Undergrad Law J. Litigation in robotic surgery: product liability or medical malpractice [Internet]. 2018. Available from: http://www.pulj.org/the-roundtable/litigation-robotic-surgery-product-liability-or-medical-malpractice [63].

64. Busch F, Kather JN, Johner C, Moser M, Truhn D, Adams LC, et al. Navigating the European Union Artificial Intelligence Act for Healthcare. NPJ Digit Med. 2024; 7:210. doi:10.1038/s41746-024-01213-6. PMID 39134637.

65. Park HJ. Patient perspectives on informed consent for medical AI: A web-based experiment. Digit Health. 2024; 10:20552076241247938. doi:10.1177/20552076241247938. PMID 38698829.

66. Gichoya JW, Banerjee I, Bhimireddy AR, Burns JL, Celi LA, Chen LC, et al. AI recognition of patient race in medical imaging: a modelling study. Lancet Digit Health. 2022; 4(6):e406-e414. doi:10.1016/S2589-7500(22)00063-2. PMID 35568690.

67. Hales M. AI in healthcare brings HIPAA privacy risks. The HIPAA E-Tool [Internet]. 2026 [cited 2026 Jul 5]. Available from: https://thehipaaetool.com/ai-in-healthcare-hipaa-privacy-and-safety-risks.

68. American College of Radiology. *ARCH-AI:* Setting the Standard for Responsible Artificial Intelligence [Internet]. Reston (VA): ACR News and Publications; [cited 2026 Jul 5]. Available from: https://www.acr.org/News-and-Publications/Responsible-Artificial-Intelligence-in-Radiology.

69. Huffman EM, Rosen SA, Levy JS, Martino MA, Stefanidis D. Are current credentialing requirements for robotic surgery adequate to ensure surgeon proficiency? Surg Endosc. 2021; 35:2104-2109. doi:10.1007/s00464-020-07608-2.

70. Proprio. Proprio receives FDA clearance for AI surgical guidance platform [Internet]. 2025. Available from: https://www.propriovision.com/news/proprio-receives-fda-clearance-for-the-worlds-first-ai-surgical-guidance-platform.

71. Goyal A, Mendoza M, Munoz AE, Macias CA, Abou-Mrad A, Marano L, et al. Artificial intelligence for real-time surgical phase recognition in minimal invasive inguinal hernia repair: a systematic review on behalf of TROGSS - the robotic global surgical society. Artif Intell Surg. 2025; 5:108. doi:10.20517/ais.2024.108.

© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

6
Wearables and Wellness: Preventing Disease Before the ER

The thing about running an emergency room is that you become an expert in what happens when the rest of the system fails.

I don’t mean that as a criticism. I mean it as a fact that shapes how you see the world. Every shift, the doors open and someone walks in with a problem that didn’t have to get this bad. The diabetic in diabetic ketoacidosis (DKA) whose glucose has been running 400 for weeks because nobody checked. The 52-year-old clutching his chest who hasn’t seen a doctor since his kids were born. The young woman with a pulmonary embolism who’s been short of breath for three days and figured it would pass.

You treat them. You stabilize them. You save the ones you can. And then the doors open again.

After enough years of this, you start seeing the pattern behind the pattern. The emergency department isn’t just where sick people go. It’s where the absence of prevention lands. Every DKA, every missed atrial fibrillation (AFib), every cancer caught six months too late is a story about a body that was sending signals no one was listening to.

I remember a night shift, when a retired teacher came in with a heart rate of 160. She’d been having palpitations for months. Months. She told her primary care doctor, who ordered a standard EKG during a 15-minute office visit. Normal sinus rhythm. Case closed. But atrial fibrillation doesn’t perform on command. It comes and goes. It hides between appointments. And by the time it showed up on my monitor, she’d already thrown a small clot. Minor stroke. Left-side weakness. She was lucky. Plenty aren’t.

I stood in the hallway after we stabilized her and thought: if she’d had a wearable that recorded her heart rhythm continuously, the AFib would have been caught in week one, not month six. No ER visit. No stroke. No 2 a.m. hallway reckoning for either of us.

That’s not a hypothetical anymore. That’s what this chapter is about: the technology to catch disease before it becomes an emergency already exists. Continuous glucose monitors, smartwatches with AFib detection, remote patient monitoring platforms that track vitals in real time. The question is no longer whether these tools work. The evidence says they do. The question is whether the healthcare system can reorganize itself around prevention rather than rescue, and whether equity concerns will be addressed before the technology widens the same gaps it has the potential to close.

I’ve spent my career on the rescue side. This chapter is about what happens when we get serious about the other side.

Executive Summary

The wrist has become medicine’s new frontier. In January 2025, Austin, Texas, man Mike Gomez was tracking workouts when his Apple Watch sent five escalating alerts. The last read: “We have detected atrial fibrillation. Contact your health professional immediately.” Gomez felt fine. He went to the hospital anyway. Doctors confirmed serious atrial fibrillation within minutes. Had he waited, his physicians told him, he would likely have suffered a life-ending stroke [1]. Gomez’s story is no longer exceptional. It’s becoming routine. Wearable devices are quietly detecting conditions that would otherwise remain silent until they became catastrophic.

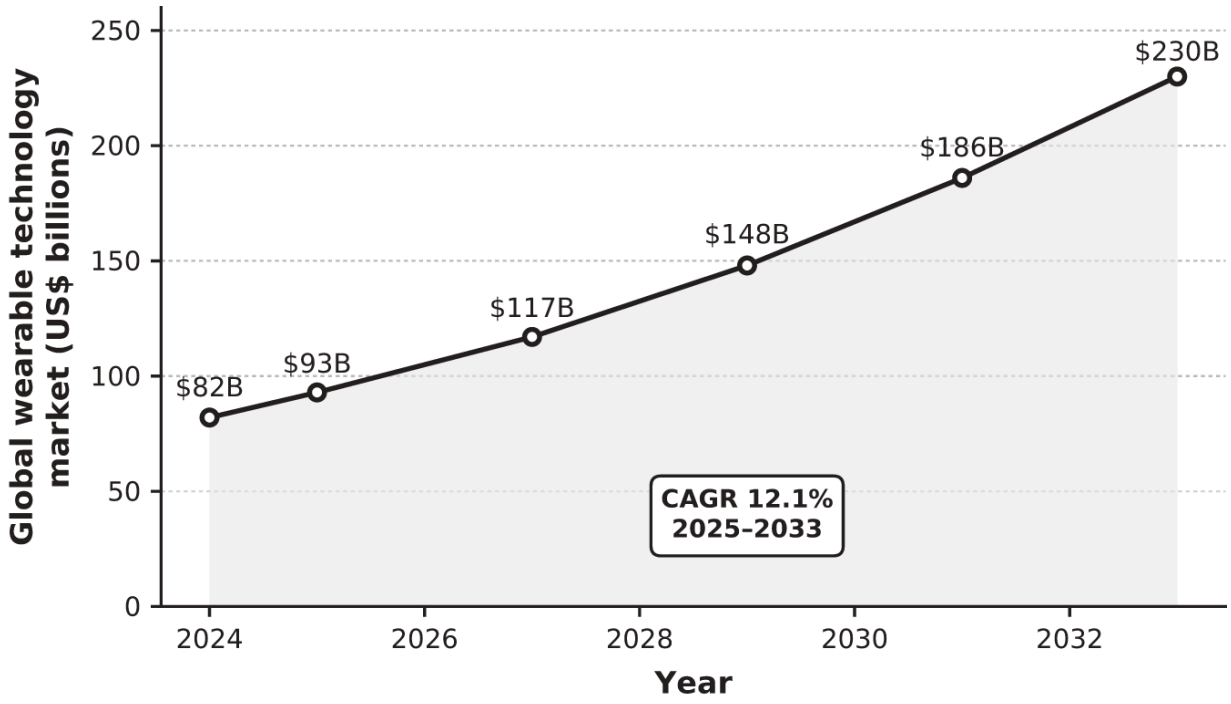
The wearable revolution, powered by artificial intelligence, is reshaping the oldest problem in medicine: finding disease early enough to stop it. The scale is immense. The global wearable technology market was estimated at $92.90 billion in 2025, projected to reach $229.9 billion by 2033 at a compound annual growth rate (CAGR) of 12.1 percent [2]. More than 62 million Americans used a health-monitoring wearable in 2024, projected to exceed 92 million by 2029 [3]. The US remote patient monitoring market was valued at approximately $14.3 billion in 2024, projected to reach $29+ billion by 2030 [4].

The clinical evidence has matured from promising to compelling. The landmark EQUAL trial showed that smartwatch-based screening detected new-onset atrial fibrillation at 4.40 times the rate of standard care, with 57

percent of detected cases completely asymptomatic [5]. The Apple Heart Study enrolled 419,297 participants and showed an 84 percent positive predictive value for AFib detection [6]. A 2025 NIH-data study of more than 14,000 participants found that an AI model trained on Fitbit data predicted hospitalizations with 91 percent accuracy [7].

Beyond cardiology, continuous glucose monitors cleared for over-the-counter purchase have placed metabolic monitoring in the hands of the more than 150 million Americans with Type 2 diabetes or prediabetes [8]. Samsung received the first FDA De Novo authorization for smartwatch-based sleep apnea detection in February 2024. Google earned FDA clearance for a feature that detects when a user's heart stops and automatically contacts emergency services [9,10].

The revolution is imperfect. Approximately 97 percent of wearable health data goes unused in clinical settings [11]. The populations most likely to benefit are the least likely to own devices. HIPAA does not protect most consumer wearable data. False positives cause psychological harm. The regulatory field is still catching up. This chapter examines both the extraordinary promise and real limitations of wearable health technology, drawing on landmark clinical trials, real-world case studies, and the latest regulatory guidance to give healthcare leaders a framework for deploying these tools responsibly (see Figure 6.1).



Source: Grand View Research/industry forecasts cited in Chapter 6.

Figure 6.1 Manual workflow versus AI-simplified workflow

Current Evidence

This section examines current evidence for remote monitoring technologies including continuous glucose monitors and photoplethysmography devices, with analysis of clinical outcomes and regulatory considerations.

The Wearable Technology Landscape

The modern wearable health ecosystem spans devices worn on the wrist, finger, arm, and torso, each with distinct sensing modalities and clinical applications. Apple Watch holds multiple FDA clearances for medical-grade features. The ECG App, available since Apple Watch Series 4 in 2018, generates a reading comparable to a single-lead electrocardiogram and can flag signs of atrial fibrillation [12]. It was the first consumer wearable to receive FDA De Novo classification for ECG capability, establishing the product code QDA under 21 CFR (Code of Federal Regulations) 870.2345 [13]. Apple Watch also uses continuous photoplethysmography to passively monitor for irregular heart rhythms suggestive of AFib, a feature that

received its own De Novo classification, product code QDB, in September 2018 [13].

On September 13, 2024, Apple received FDA 510(k) clearance for a sleep apnea notification feature on Apple Watch Series 9, Series 10, and Apple Watch Ultra 2, using the watch's accelerometer to detect subtle wrist movements associated with breathing disruptions during sleep [14]. The feature analyzes data collected over at least 10 nights within a 30-day period and alerts users to possible moderate-to-severe obstructive sleep apnea risk. Apple Watch Series 11 introduced the Hypertension Notification Feature (HTNF), cleared under 510(k) K250507 in September 2025 for hypertension-related notifications [13].

Samsung received FDA De Novo authorization on February 9, 2024, for sleep apnea detection in the Samsung Health Monitor app, the first smartwatch-based sleep apnea feature authorized by the FDA, predating Apple's clearance by several months [15]. Clinical validation data showed 82.7 percent sensitivity and 87.7 percent specificity in a 470-subject study [15]. Samsung's Irregular Heart Rhythm Notification feature was cleared in May 2023, and its ECG app received a further 510(k) clearance in August 2024 [13].

Google's Pixel Watch 4, announced in August 2025 and released in October 2025, features ECG assessment, irregular heart rhythm notifications, advanced sleep tracking, and a personalized AI health coaching capability powered by Google Gemini [16]. On February 27, 2025, the FDA cleared Google's Loss of Pulse Detection feature, a first-of-its-kind capability that detects when a user's heart stops beating and automatically contacts emergency services if the user is unresponsive [10].

The Oura Ring (Gen 3 and Gen 4) delivers continuous tracking of heart rate variability, resting heart rate, respiratory rate, blood oxygen during sleep, skin temperature, activity, and sleep stages. Published accuracy data shows heart rate accuracy of 99 percent and HRV accuracy of 98 percent compared to ECG reference standards, and sleep tracking accuracy of 79 percent agreement with clinical polysomnography [17]. A 2025 systematic review in OTO Open concluded that the Oura Ring shows "strong agreement with medical-grade sleep studies across multiple sleep parameters" [18].

AI Algorithms Powering Wearable Health Monitoring

What sets modern wearable health technology apart from earlier generations of fitness trackers is the sophistication of the AI processing sensor data. Raw physiological signals, optical waveforms, accelerometer readings, temperature changes, become clinically useful only after algorithmic transformation. At the foundation lies photoplethysmography, the optical method by which nearly all consumer wearables measure heart rate and blood oxygen: LEDs light up the skin, and photodetectors measure changes in light absorption as blood volume pulses with each heartbeat. Machine learning algorithms convert these waveforms into heart rate, heart rate variability, and respiratory rate measurements, and detect irregular rhythm patterns indicative of AFib.

Advanced multisensor fusion approaches combine PPG (photoplethysmography), ECG, accelerometry, skin temperature, and blood oxygen signals through ensemble models to achieve meaningfully better classification accuracy for complex conditions like sleep staging and exertion responses [19]. Foundation models trained on more than 2.5 billion hours of wearable data from 162,000 individuals have shown that behavioral metrics, sleep consistency, activity rhythms, mobility trends, often predict health outcomes more effectively than raw sensor data alone [20]. On-device AI processing is increasingly common, with sensitive health computations performed locally to protect privacy. Apple's sleep apnea detection algorithm runs entirely on the device and paired iPhone rather than on cloud servers [14].

Clinical Validation Evidence

The Apple Heart Study, published in the *New England Journal of Medicine* in November 2019, enrolled 419,297 participants across all 50 states over eight months [6]. Only 0.52 percent of participants received irregular pulse notifications, a reassuringly low false positive rate for a population-scale screening tool. Among 450 participants who returned analyzable ECG patches, 34 percent had confirmed atrial fibrillation. The positive predictive value was 0.84, meaning 84 percent of smartwatch irregular-pulse alerts were concordant with confirmed AFib on simultaneous ECG patch recording. Among those who received notifications and completed a 90-day

follow-up survey, 57 percent contacted healthcare providers outside the study [6].

The Fitbit Heart Study enrolled 455,699 US adults aged 22 or older [21]. Among 1,057 participants who received notifications and returned analyzable ECG patches, 32.2 percent had confirmed atrial fibrillation. The primary positive predictive value of the algorithm, measuring concurrent AFib at the time of a new notification, was 98.2 percent. Most detected AFib episodes occurred during sleep and were completely asymptomatic [21].

The most rigorous evidence to date comes from the EQUAL trial, published in the *Journal of the American College of Cardiology* on January 22, 2026 [5]. This prospective, multicenter, randomized controlled trial enrolled 437 patients aged 65 or older (median age 75) with elevated stroke risk and no prior history of atrial fibrillation. Participants were randomized to receive an Apple Watch for six months of continuous monitoring plus telehealth cardiac review, or to standard care. New-onset AFib was diagnosed in 21 patients (9.6 percent) in the smartwatch group compared to only 5 patients (2.3 percent) in the standard care group, a hazard ratio of 4.40. The number needed to screen (NNS—the number of individuals who must be screened to identify one true positive) was 14 patients to detect one additional AFib case. Fifty-seven percent of AFib cases in the smartwatch arm were asymptomatic. All patients diagnosed were immediately started on oral anticoagulation. Commenting on the findings, University of Utah cardiologist T. Jared Bunch, MD, observed: “This trial showed that we can’t rely on symptoms alone” [5].

A meta-analysis of 11 studies comprising 4,241 participants, published in *JACC: Advances* in January 2025, found that the Apple Watch ECG achieved a pooled sensitivity of 94.8 percent and specificity of 95 percent for detecting atrial fibrillation, with an area under the receiver operating characteristic curve of 0.96 [22].

I’ve seen this exact dynamic play out in my ER: the patient who comes in because their watch alarmed is usually ahead of the patient who waited for symptoms. The EQUAL trial just put a number on what emergency physicians have been observing for years.

CGM Evidence and Outcomes

The global continuous glucose monitor (CGM) market was valued at approximately $5.9 billion in 2024, projected to reach $13.6 billion by 2034 [23]. Abbott's FreeStyle Libre system is used by approximately 7 million people globally, with Abbott generating $6.8 billion in Diabetes Care revenue in 2024 growing nearly 20 percent organically [24]. On March 5, 2024, Dexcom received FDA clearance for Stelo, the first over-the-counter continuous glucose biosensor in the United States, available without a prescription for adults with Type 2 diabetes not using insulin [25].

A 2023 systematic review in the Journal of Clinical Endocrinology and Metabolism analyzed 14 randomized controlled trials and found a pooled HbA1c reduction of 0.32 percent compared to standard blood glucose monitoring [26]. The Mobile trial, a landmark multicenter RCT of 175 adults with Type 2 diabetes on basal insulin, found CGM use reduced HbA1c from 9.1 percent to 8.0 percent at 8 months versus 9.0 to 8.4 percent with finger-stick monitoring and gave CGM patients an additional 3.6 hours per day of time in the target glucose range [27]. For Type 1 diabetes, a 2025 review found HbA1c improvements of 0.4 to 0.6 percent with CGM (continuous glucose monitor), a 72 percent reduction in hypoglycemic events below 54 mg/dL in one study, and a fourfold decrease in severe hypoglycemia-related hospitalizations in retrospective data [28].

Wearable SMART Patches and Emerging Biosensors

Beyond smartwatches and rings, smart biosensor patches represent a rapidly growing category. The global wearable patch market was estimated at $9.95 billion in 2024, projected to reach $15.71 billion by 2030 [29]. ECG patches such as the iRhythm Zio enable extended cardiac rhythm monitoring up to 14 or more days, while advanced patches can detect sweat biomarkers including electrolytes, cortisol, glucose, and lactate. A 2026 review in *BioMedical Engineering OnLine* described smart patches as playing a complex role encompassing continuous health monitoring, controlled drug delivery, and tissue regeneration support [30]. The VitalPatch and ReDS Wearable System both hold FDA approval specifically for heart failure patient monitoring.

Real-World Case Studies

These cases illustrate practical implementation challenges and successes in real-world settings, demonstrating integration of remote monitoring technologies across different healthcare systems.

From Alert to Ambulance: Individual Stories of Wearable Life Saves

No statistic communicates the clinical value of wearable health monitoring as directly as the stories of patients whose lives it has saved. Frank Haggerty, 73, a retired accountant from Lawrenceville, New Jersey, was asleep in February 2024 when his Apple Watch SE triggered an alert showing his heart rate at approximately 30 beats per minute, far below the normal range of 60 to 100. Despite feeling completely fine, he was rushed by ambulance to the hospital. Doctors diagnosed him with complete heart block, a malfunction of the heart's electrical system requiring immediate pacemaker implantation. His cardiologist, Dr. Keith Wolfson of Capital Health, stated without equivocation that Haggerty would not have survived without the watch's alert. "Quite frankly, had I not had the watch on, I wouldn't be sitting here today," Haggerty told CBS News [31].

Kimmie Watkins, 29, of Cincinnati, Ohio, had no history of heart problems when she felt lightheaded in June 2023 and decided to rest. Her Apple Watch alerted her approximately 90 minutes later that her heart rate had been elevated above 178 beats per minute for more than 10 minutes. She went to the emergency room, where doctors discovered a saddle pulmonary embolism, a blood clot straddling the arteries feeding both lungs. University of Cincinnati cardiologist Dr. Richard Becker noted that individuals with this type of embolism face a 50 percent mortality risk. The watch alert, delivered to a 29-year-old with no obvious risk factors, quite likely saved her life [32].

In 2025, Roberto Gallart, a regular gym-goer in Rio de Janeiro, noticed sudden chest pain during a workout. He checked the ECG feature on his Samsung Galaxy Watch 6 and found four "inconclusive heart rhythm" readings from that day. He rushed to the hospital, where physicians discovered one of his main coronary arteries was completely blocked and the other two critically narrowed. He was on the verge of a myocardial

infarction. “The ECG feature was instrumental in saving my life. Without its warning, I may not have sought help in time,” Gallart told Samsung’s global newsroom [33].

The Oura Ring also showed the potential of wrist-worn biosensors for early infectious disease detection during the COVID-19 pandemic. The TemPredict Study, conducted by UCSF researchers and enrolling 63,153 Oura Ring wearers, including approximately 3,400 frontline healthcare workers, found that an algorithm combining skin temperature, heart rate, HRV, respiratory rate, and activity data detected COVID-19 onset an average of 2.75 days before participants sought conventional diagnostic testing, with an area under the receiver operating characteristic curve of 0.819 [34].

Health System Implementations at Scale

The Veterans Health Administration’s Remote Patient Monitoring–Home Telehealth program stands as one of the largest and most outcome-documented wearable devices (distinct from remote patient monitoring systems, which are clinical-grade medical devices) monitoring deployments in the world. Serving more than 132,000 Veterans per year and generating over 7.7 million telehealth episodes in fiscal year 2025, a 12 percent increase over the prior year, the program remotely collects biometric data including blood pressure, weight, and blood glucose from veterans at home [35]. Veterans enrolled show a 41 percent reduction in hospital admissions and approximately a 70 percent reduction in hospital bed days of care. In one monitored subnetwork during fiscal year 2024, the program served 6,196 unique veterans, reduced bed days of care by 50 percent, and achieved a 94 percent veteran satisfaction rate [35].

The real human payoff is visible in individual incidents: during a routine video visit, Veterans Administration (VA) care coordinator Kenneth Horseman noticed a patient was having difficulty speaking and appeared pale, immediately directing the veteran to seek emergency care. The veteran was in the ICU the next day but recovering [35].

Mass General Brigham’s RPM hypertension and cholesterol program enrolled over 10,000 patients, analyzing nearly half a million blood pressure readings and 130,000 laboratory reports. Both blood pressure and

cholesterol levels declined significantly. Critically, the improvements were comparable across race, ethnicity, and language groups, suggesting that well-designed RPM programs need not replicate the equity disparities that characterize broader healthcare access [36].

At Northwestern Medicine, a study of an Omron VitalSight blood pressure monitoring system integrated with Epic across six Chicago-area primary care practices evaluated 288 Medicare patients randomized to the RPM program versus 1,152 propensity-matched controls. The proportion with controlled blood pressure nearly doubled in the RPM group (from 35 percent to 70–73 percent) compared to a modest increase among controls (from 39 percent to 51–59 percent). RPM patients showed a 10.4 mmHg decrease in systolic blood pressure versus 4.7 mmHg in controls, and were 1.7 to 2.5 times more likely to have controlled blood pressure at all measured time points [37].

In my experience running an ER, what these numbers represent is not just better metrics: it’s fewer patients rolling through the door in hypertensive crisis. In clinical research, a Mayo Clinic 12-month randomized trial involving 184 physicians found that those randomly assigned to wear a smartwatch and receive personal health data saw a 54 percent reduction in the overall odds of burnout compared to those without a study smartwatch, showing that wearable technology’s clinical applications extend to physician wellness, not just patient care [38].

Employer and Insurer Wellness Programs

John Hancock became the first major US life insurer to tie every policy to health tracking, offering Apple Watches to all eligible Vitality Plus members. A Rand Europe study of more than 400,000 program participants, described at publication as the world’s largest behavior-tech study based on verified data, found that participants who wore an Apple Watch averaged a 34 percent sustained increase in physical activity. In the United States specifically, high-intensity activity days increased by 52 percent, and participants with high inactivity levels and high BMI saw a 200 percent increase in activity [39].

UnitedHealthcare launched its UnitedHealthcare Rewards program in February 2023, offering eligible members up to $1,000 per year through

integration with wearable devices and smartphones, rewarding step goals, activity duration targets, sleep tracking over 14 consecutive nights, and completion of biometric screenings [40]. Oura Health announced a clinical care partnership with Maven Clinic in June 2025, integrating Oura biometric data directly into the Maven care team's clinical workflow [41]. In October 2024, Essence Healthcare became the first Medicare Advantage plan to partner with a smart ring device, providing Oura Rings at no cost to certain Medicare Advantage members starting in 2025 [42].

These programs trace a clear trajectory. Wearable devices are no longer merely consumer gadgets for the health-curious. They're becoming infrastructure for a new model of preventive care, one that monitors continuously, intervenes early, and measures outcomes at population scale.

Risks and Limitations

While remote monitoring offers significant clinical benefits, important limitations exist including device accuracy variability, user adherence challenges, and potential for false alerts requiring careful clinical interpretation.

Clinical Accuracy: The Photoplethysmography Challenge

No technology deployed at population scale is without limitations, and wearable health monitoring is no exception. The foundational sensing technology in virtually all consumer wearables, photoplethysmography, carries inherent sources of inaccuracy. Three major known sources exist: motion artifacts, skin type variation, and signal crossover. Motion artifacts occur when repetitive wrist movements cause the optical sensor to lock onto the motion frequency rather than the cardiovascular cycle. Skin tone variation affects photoplethysmography (PPG) accuracy because melanin absorbs more green light than lighter pigmentation, potentially degrading signal quality in individuals with darker skin. Signal crossover causes periodic rhythmic motion, walking or jogging, to be mistaken for the cardiovascular pulse cycle [19].

A 2020 *Nature Digital Medicine* study found that absolute heart rate error during physical activity was, on average, 30 percent higher than during rest across consumer and research-grade wearables [19]. A 2026 living systematic review of Apple Watch measurement accuracy in *npj Digital Medicine* found that while AFib detection showed acceptable specificity (0.91), its sensitivity was lower (0.79), and blood oxygen readings had a clinically significant range of limits of agreement spanning nearly 8 percentage points [43].

False Positives and the Anxiety Cascade

A 16 percent false-positive rate in a smartwatch AFib notification system may sound acceptable in clinical terms. At population scale, across tens of millions of device users, it translates to a substantial volume of unnecessary cardiology visits, anxious patients, and disrupted clinical workflows. A 2025 narrative review in the *Journal of Clinical Medicine* found that "patients encountering false-positive alerts experience heightened health anxiety, reduced perceived well-being, and compromised adherence to prescribed monitoring protocols" [44]. The paper identified premature atrial and ventricular contractions as major sources of false-positive AFib alerts because they mimic AFib-induced rhythm irregularity on PPG analysis. The *Cleveland Clinic Journal of Medicine* flagged: "The high false-positive rate for detecting atrial fibrillation in young, otherwise healthy populations is a significant limitation of consumer-based devices that may lead to increased anxiety and unnecessary healthcare use" [45].

A further accuracy concern involves the 27.9 percent of initial Apple Watch ECG recordings that were unclassifiable in one systematic review, requiring physician review to resolve. When unclassifiable recordings were treated as incorrect, sensitivity fell to 66.2 percent and specificity to 73.4 percent [46].

SpO_2 Accuracy and Skin Tone: A Documented Safety Disparity

Among the most significant accuracy concerns is the well-documented inaccuracy of pulse oximetry in individuals with darker skin pigmentation. The FDA has formally acknowledged that SpO_2 is frequently overestimated in individuals with darker skin compared to gold-standard arterial blood gas

measurements, meaning darker-skinned patients may have clinically significant oxygen desaturation that pulse oximeters fail to detect. The FDA noted this “occult hypoxemia” may “lead to delayed medical attention” and “adverse health outcomes” [47]. In January 2025, the FDA published draft guidance proposing stricter premarket requirements for pulse oximeters, including mandatory testing across diverse skin pigmentations using the Monk Skin Tone Scale [48]. The December 2024 FDA guidance on predetermined change control plans explicitly recommended that manufacturers consider “the unique characteristics of their device’s intended use populations (such as ethnicity, gender, and disease severity)” in planning algorithm updates [49]. Inaccurate SpO_2 readings feed into hospital deterioration algorithms that guide supplemental oxygen decisions and ICU triage, potentially amplifying existing racial disparities in care at the most critical moments.

Calorie Estimation: The Weakest Link

The least accurate metric across consumer wearable devices is energy expenditure estimation. A meta-analysis of Apple Watch models found mean absolute percentage errors of 4.43 percent for heart rate and 8.17 percent for step counts, but 27.96 percent for energy expenditure, an inaccuracy observed across all types of users and activity types [50]. Fitbit energy expenditure accuracy is similarly limited, with one analysis finding an average error of approximately 34.4 percent [50]. For clinical programs relying on wearable-derived caloric data, this limitation demands explicit acknowledgment.

Psychological Risks: Orthosomnia, Cyberchondria, and Alert Fatigue

Orthosomnia, excessive preoccupation with achieving perfect sleep scores, has emerged as a recognized phenomenon among wearable sleep tracker users. A 2025 *JMIR* study in a Canadian sample of 1,153 participants found that wearable users reported 13 minutes longer sleep onset latency and approximately one hour less sleep per night compared to nonusers, and endorsed more severe insomnia symptoms [51]. Cyberchondria, health anxiety fueled by misinterpretation of continuous digital self-tracking data, represents a related risk. A 2025 *Nature npj Digital Medicine* analysis

documented that “some alerts or abnormal readings from wearables can serve as analogous ‘false alarms,’ where benign variations trigger unnecessary concern. In the digital health context, this issue potentially operates at a larger scale and frequency, with consumer-oriented tools continuously generating metrics and alerts that often involve interpretation without clinical oversight” [52].

Alert fatigue is a well-documented patient safety problem. Research has shown that 80 to 99 percent of ECG monitor alarms in clinical environments are false or clinically insignificant [53]. The Cleveland Clinic’s Cosmos study found that with default alarm thresholds, 6 percent of patients activated more than 24 alarms per day, with six outlier patients accounting for 46 percent of total alarms [53]. Adding consumer RPM alerts without intelligent filtering risks further desensitizing care teams to genuinely important signals.

Data Going Unused

The most troubling limitation of the current wearable ecosystem is not what devices can’t measure, but what happens to the data they generate. Approximately 97 percent of wearable health data goes unused in clinical settings [11]. This is not primarily a technology problem. It’s a workflow, integration, and clinical protocol problem. In the absence of standardized fast healthcare interoperability resources (FHIR) integration, structured care pathways for acting on wearable alerts, and clear protocols for incorporating continuous biometric data into clinical decision-making, the most sophisticated sensors in history generate noise rather than signal.

In my experience, the data going unused problem is not technical; it’s organizational. Every ER I’ve worked in has had data flowing past clinicians without ever reaching them in actionable form. Wearables just add another stream.

Regulatory and Legal Landscape

Remote monitoring technologies are subject to evolving regulatory frameworks across jurisdictions, including FDA oversight of medical

devices, data privacy requirements under HIPAA and international standards, and reimbursement considerations affecting clinical adoption.

The FDA's Tiered Approach to Wearable Health Devices

The FDA regulates wearable health devices under the Federal Food, Drug, and Cosmetic Act when a device is "intended for use in the diagnosis of disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease." The agency's approach distinguishes three tiers [13]. Non-device software intended solely for general wellness, step counters and calorie trackers, is excluded from the device definition under the 21st Century Cures Act of 2016. General wellness devices presenting low risk are subject to enforcement discretion. Wearables making diagnostic, therapeutic, or disease-monitoring claims must clear or obtain approval through the 510(k), De Novo, or PMA pathways.

The September 2018 dual De Novo grants for Apple Watch were landmark events in wearable regulation history. Apple Watch Series 4 became the first large-scale consumer wearable product that was itself an FDA-cleared medical device [13]. The ECG App received De Novo classification Den180044, establishing product code QDA ("Electrocardiograph software for over-the-counter use") under 21 CFR 870.2345. The Irregular Rhythm Notification Feature received simultaneous De Novo classification Den180042, establishing product codeQDB ("Photoplethysmography analysis software for over-the-counter use") under 21 CFR 870.2790. These classifications created the predicate device framework that subsequent wearable manufacturers, including Fitbit, Samsung, and Google, relied upon for their own 510(k) clearances [13]. Samsung's sleep apnea detection received De Novo authorization DEN220033 in February 2024, simultaneously establishing a new device type and the predicate that Apple used seven months later for its own sleep apnea clearance via 510(k) K240929 [14,15]. Google's Loss of Pulse Detection received FDA clearance in February 2025 as another novel first-of-its-kind wearable feature [10].

The Predetermined Change Control Plan: Regulating Adaptive AI

Among the most consequential regulatory developments for AI-powered wearable health software is the FDA's December 2024 final guidance on "Predetermined Change Control Plans for AI-Enabled Devices" [49]. A PCCP allows manufacturers of AI-enabled devices with existing marketing authorization to implement future device modifications, including retraining machine learning algorithms on new real-world data, without submitting new marketing applications, provided the modifications fall within a pre-approved protocol. This framework addresses a fundamental mismatch between the continuous improvement cycle of machine learning systems and the historically static nature of FDA device clearances. Apple's irregular rhythm notification feature was among the first wearable applications to include a PCCP as part of its 510(k) K231173 clearance in July 2023, specifically authorizing future algorithm retraining within a pre-defined performance envelope [13]. The December 2024 final guidance expanded the PCCP's scope from machine learning-only to all AI-enabled devices and explicitly added recommendations for manufacturers to consider diverse populations, including different ethnicities, gender groups, and disease severity, in change control planning [49].

The Software as a Medical Device Framework

The FDA's Software as a Medical Device (SaMD) framework defines SaMD as "software intended to be used for one or more medical purposes that perform these purposes without being part of a hardware medical device." By July 2025, the FDA had authorized more than 1,250 AI- or ML-enabled devices under this framework, with wearable-adjacent applications representing a growing share [54]. The International Medical Device Regulators Forum's SaMD guidance, which the FDA has adopted, classifies SaMD on a four-tier scale based on the significance of the information provided by the software and the healthcare situation or condition. Consumer wearable apps with medical claims typically fall at SaMD Categories I or II under this classification [54].

In the United Kingdom, the Medicines and Healthcare products Regulatory Agency defines "medical device software" under the Medical Devices

Regulations 2002 as amended, and wearable health features with clinical claims require United Kingdom Conformity Assessed (UKCA) marking post-Brexit, with Medicines and Healthcare Regulatory Agency (MHRA) guidance specifically addressing software of unknown provenance [55]. European CE marking requirements under the Medical Device Regulation (EU 2017/745) impose more stringent clinical evidence requirements than FDA 510(k) for equivalent product categories, creating regulatory asymmetry that complicates global product strategy for wearable health manufacturers [55].

HIPAA and the Wearable Data Gap

One of the most consequential and least understood legal realities of consumer wearable health technology is that most of it falls outside HIPAA's protections. HIPAA applies to covered entities, primarily healthcare providers, health plans, and healthcare clearinghouses, and their business associates. Consumer wearable companies are not covered entities. Apple, Fitbit, Garmin, Samsung, and similar manufacturers are not subject to HIPAA restrictions on how they collect, share, or sell consumer health data generated by their devices, unless that data is transmitted to a covered entity as part of a defined clinical workflow [56].

In practice, heart rate data, sleep patterns, menstrual cycles, blood oxygen readings, and activity logs collected by consumer wearables can be shared with third parties, including advertisers, data brokers, and insurers, without explicit user consent in most US states. The FTC Act's prohibition on deceptive practices provides some floor-level protection, and the FTC has brought enforcement actions against health app companies that shared sensitive data contrary to their privacy policies [57]. The FTC's 2021 enforcement action against Flo Health found that the period-tracking app shared health data, including pregnancy status and health conditions, with Facebook, Google, and other third parties [58–62].

In July 2024, the FTC expanded its Health Breach Notification Rule to cover health apps and wellness platforms "associated with wearable technologies," requiring notification within 60 days of breach discovery and extending coverage to unauthorized disclosures, not just hacking incidents, with civil penalties reaching $43,792 per violation per day [63]. Washington's My Health My Data Act, effective March 2024, provides the

broadest state-level coverage: requiring opt-in consent for collection of any "consumer health data," expressly including wearable-derived metrics, with a private right of action for violations [64]. As of early 2026, proposed federal legislation, the Health Information Privacy Reform Act, would extend federal privacy baselines to wearables, but the bill remained in committee [65].

The Digital Divide: Wearable Health for Whom?

One of the most consequential ethical challenges is the stark stratification of access along lines of income, race, age, and geography. A nationally representative study of 23,974 US respondents found that adults aged 65 or older had 82 percent lower odds of owning a wearable compared to 18- to 24-year-olds, rural residents had 35 percent lower odds, and uninsured individuals had 59 percent lower odds [66]. Those earning above $200,000 per year had 2.27 times higher odds of ownership compared to those earning below $50,000 [66]. A 2025 study in the *Journal of Patient Experience* found that African American and Latino patients were 17 percent and 15 percent less likely, respectively, to use wearable devices. Black patients were also under-enrolled and under-referred for remote patient monitoring programs for heart failure [67,68].

The equity paradox is acute. The populations most likely to benefit from preventive health monitoring, those with higher burdens of cardiovascular disease, diabetes, and hypertension, are precisely those with the lowest device access. In the US alone, excess healthcare costs due to racial health disparities are estimated at $93 billion annually [69]. If wearable-based preventive care is not deployed equitably, it risks widening that gap rather than narrowing it.

Employer-Mandated Wearables and EEOC Guidance

The December 2024 Equal Employment Opportunity Commission (EEOC) fact sheet on employer-furnished wearable devices made several important determinations [70]. An employer's wellness program that collects physical or mental health condition data via a wearable device may constitute a prohibited medical examination under the Americans with Disabilities Act. Mandatory wearable use does not satisfy the ADA's voluntary wellness program requirement. The EEOC directly identified algorithmic bias as a

compliance concern: biometric devices that fail to calibrate for darker skin tones or larger body sizes may produce skewed results, creating discriminatory employment outcomes [70].

Insurance Incentives and the Ethics of Continuous Monitoring

The American Medical Association has published substantive analysis of ethical concerns surrounding insurers' use of wearable data. While rewards programs currently use wearable data to incentivize healthy behavior, the AMA identified "no regulations stopping them from doing the opposite, penalizing policyholders with higher premiums or denying coverage based on activity or health data" [71]. The AMA also identified risks that inaccurate device readings could result in lower rewards for genuinely healthy behaviors, or trigger false positives for conditions like AFib that lead to unnecessary treatment [71]. A Brooklyn Law Review analysis notes: "There are no clear regulatory controls preventing insurers from using wearable data to influence premium pricing or coverage decisions, despite the significant personal health information at stake" [72].

Leadership Implementation Framework

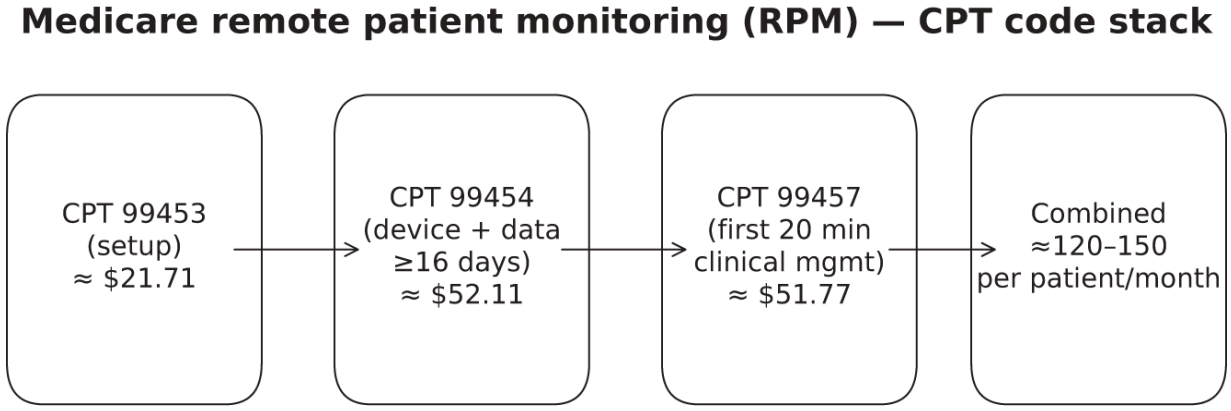
For healthcare CEOs, CIOs, CMOs, and system leaders, the question is no longer whether wearable health technology deserves a place in clinical strategy. The evidence is clear that it does. The question is how to deploy it responsibly, sustainably, and at scale.

RPM Program Deployment and Reimbursement

Medicare reimbursement for remote patient monitoring has matured into a financially viable framework for health systems. The current procedural terminology (CPT) code structure includes CPT 99453 for one-time device setup and patient education, reimbursing approximately $21.71; CPT 99454 for device supply and data transmission per 30-day period requiring at least 16 days of readings, reimbursing approximately $52.11; and CPT 99457 for the first 20 minutes of clinical management per month, reimbursing approximately $51.77 [73]. Combined, Medicare RPM programs can generate approximately $120 to $150 per patient per month, exceeding

$1,000 per patient annually [73]. Medicare payments for RPM services reached $536 million in 2024, a 31 percent increase from 2023, with nearly 1 million Medicare enrollees receiving RPM services, a 27 percent year-over-year increase [74]. Forty-two state Medicaid programs now cover RPM for certain populations, and by 2023, 81 percent of clinicians reported using RPM, a 305 percent increase since 2021 [74] (see Figure 6.2).

Program design matters as much as program existence. The VA's RPM-HT program, the Northwestern Medicine hypertension pilot, and UPMC's high-risk patient program, which achieved a 76 percent lower 90-day readmission rate with patient satisfaction exceeding 90 percent, all share common design elements: EHR integration, structured clinical escalation protocols, dedicated care coordination teams, and patient-centered engagement strategies [75,76].



Source: CMS Medicare Physician Fee Schedule (manuscript Chapter 6).

Figure 6.2 AI-automated billing flow: encounter to payment

FHIR Integration and the Interoperability Challenge

The most significant technical barrier to scaling wearable-based care is interoperability. Consumer wearables operate in proprietary ecosystems, Apple HealthKit, Fitbit/Google, Garmin Connect, WHOOP, Samsung Health, that resist integration with electronic health record platforms [11]. Data formats are proprietary, metrics are inconsistently defined, and no universal standard yet maps consumer-grade biometric data to clinical data structures. The HL7 Fast Healthcare Interoperability Resources standard has emerged as the leading technical framework for mapping wearable data

to EHR-compatible formats. Implementation requires custom API configuration, dedicated integration infrastructure, and expert data transformation. It's not a plug-and-play solution [76].

Healthcare CIOs must budget for this integration work explicitly and evaluate vendors on the maturity of their FHIR implementation. The most successful large-scale RPM programs, Northwestern Medicine's Omron VitalSight integration with Epic and Mass General Brigham's continuous biometric monitoring program, achieved clinical impact precisely because data flowed into existing clinical workflows without friction [37,76].

Vendor Evaluation Criteria

When evaluating wearable health technology vendors, healthcare executives should apply a structured assessment framework addressing five domains.

FDA regulatory status: Is the specific health feature FDA-cleared, and under what pathway? De Novo and 510(k) clearances confer different levels of predicate support. General wellness devices may be appropriate for certain applications but should not be deployed for clinical decision support.

Clinical validation evidence: Are published peer-reviewed studies available in populations comparable to your patient population? Validation studies in young, healthy populations may not generalize to older, higher-acuity clinical populations. Ask vendors specifically about accuracy data stratified by age, race, and comorbidity burden.

Data security and privacy architecture: Does the vendor offer business associate agreements for clinical deployments? What are the vendor's third-party data sharing practices? Does data remain on-device or is it transmitted to cloud platforms?

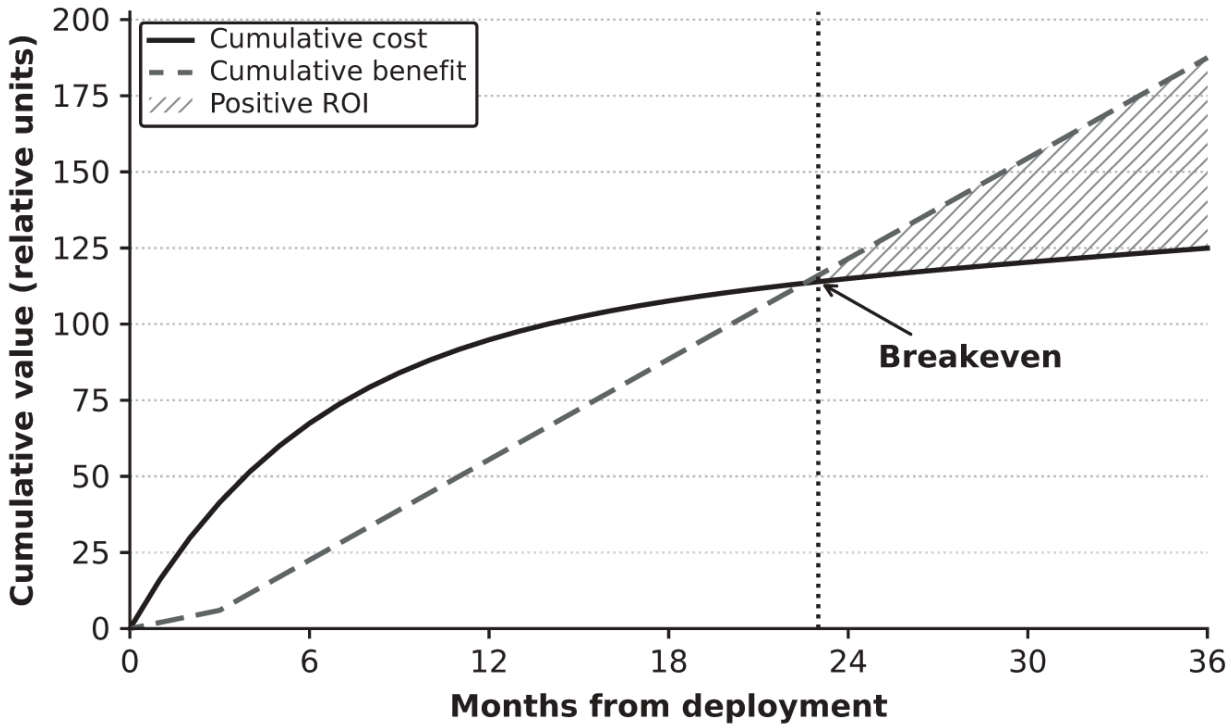
EHR integration maturity: Can the vendor show actual FHIR integration with your EHR platform, not just API documentation?

Equity and accessibility: What is the device's accuracy profile across different skin tones? Is it accessible to patients with limited digital literacy? Does the vendor offer subsidized access for low-income populations?

Change Management and Clinician Buy-In

Technology adoption in healthcare fails far more often from change management failure than from technical inadequacy. Clinician resistance to wearable data stems from genuine concerns about workflow burden, liability exposure, and data quality. A qualitative study of secondary care providers found that adoption was hindered by provider perceptions of information overload, concerns about liability from lack of timely review, and fears about patient privacy [77]. Addressing these concerns requires structured education programs explaining device limitations, clear clinical protocols specifying which data triggers escalation and who is responsible for acting on it, and intelligent filtering systems that surface actionable signals while suppressing noise.

The economic case for change is compelling. RPM programs have shown 38 percent average reductions in hospital admissions and 51 percent reductions in emergency department visits [78]. Geisinger Health Plan reported a 77 percent reduction in all-cause readmission risk for Medicare Advantage members in its heart failure RPM program [79]. A systematic review in Mayo Clinic Proceedings: Digital Health found wearable health technologies can increase quality-adjusted life years, with cost-effectiveness probability ranging from 66 to 100 percent across analyzed scenarios [80]. Wearable cardiovascular monitoring was associated with a 43 percent lower likelihood of hospital readmission among coronary artery disease patients, with cost savings of approximately $6,000 per patient [81]. Global cost savings from wearable technology in healthcare are projected at approximately $200 billion [82] (see Figure 6.3).



Conceptual ROI curve. Magnitudes are illustrative; actual recovery depends on use case, scale, and reimbursement mix (Chapter 6).

Figure 6.3 Institutional ROI from AI implementation: time and cost recovery

Future Outlook

Future developments in remote monitoring include advancing artificial intelligence for predictive analytics, expanding interoperability standards, improving user interface design, and addressing equity in access to technologies across diverse populations.

Next-Generation Wearable Capabilities

The wearable health technology field of 2026 is a foundation, not a ceiling. Several next-generation capabilities are advancing rapidly from research to clinical application. Blood pressure monitoring without a cuff represents one of the most significant unsolved problems in continuous wearable health data. Apple's Watch Series 11 introduced blood pressure trend notifications cleared by the FDA in September 2025 for hypertension-related alerts [13]. Oura is pursuing FDA approval for a blood pressure

monitoring feature as one of its most ambitious next-generation initiatives [83]. The American Heart Association's December 2025 scientific statement was clear: cuffless blood pressure devices "show great promise…however, they are not yet proven accurate enough to diagnose high blood pressure or guide treatment decisions," and the 2025 AHA/ACC High Blood Pressure Guideline explicitly recommended against using cuffless devices for clinical diagnosis or treatment [84]. Solving this accuracy challenge, particularly for nighttime blood pressure assessment where cardiovascular risk data is most informative, would reshape the clinical utility of wearable devices.

Samsung announced in October 2025 that Galaxy Watches will gain the ability to detect left ventricular systolic dysfunction, a weakening of the heart muscles responsible for approximately half of all heart failures worldwide. The feature had already received regulatory approval in South Korea, with clearance in other markets pending [85]. When broadly available, this capability would extend wearable cardiac monitoring beyond rhythm detection into structural cardiac disease assessment.

AI-enhanced CGMs are advancing from glucose measurement to glucose prediction. A meta-review by University at Buffalo researchers published in *npj Digital Medicine* in November 2025, analyzing 60 peer-reviewed studies, found that AI-enhanced CGMs can predict glucose changes one to two hours in advance, turning CGMs from retrospective measurement tools into prospective warning systems. The researchers described this as "turning CGMs from a rear-view mirror into a heads-up display" [86].

The Regulatory Horizon

The FDA's approach to AI-powered wearable devices is evolving toward a lifecycle model rather than a static premarket clearance paradigm. The PCCP framework enables manufacturers to pre-authorize algorithm improvements, reducing the regulatory friction that would otherwise slow beneficial AI updates [49]. The January 2025 draft guidance called "AI-Enabled Device Software Functions Lifecycle Management" signals that the FDA expects continuous performance monitoring tied to algorithm changes and labeling updates when modifications are implemented [54].

The Tempo pilot, launched in February 2026 in partnership with CMS, tests a new model for promoting access to certain digital health devices while maintaining safety, potentially accelerating the pathway from FDA clearance to Medicare coverage [13]. The World Health Organization's updated "Global Strategy on Digital Health 2020–2027," released in December 2025, reports that at least 129 countries have established national digital health strategies, creating a global policy environment that, while fragmented, is increasingly receptive to wearable-based preventive care at population scale [87].

The Convergence Opportunity

The most significant near-term opportunity lies at the intersection of three trends: the maturation of wearable biosensors, the emergence of AI foundation models trained on longitudinal biometric data, and the expansion of telehealth infrastructure. The Heartline Study, a large randomized trial of more than 34,000 older adults examining how an Apple Watch-based app program affects early AFib diagnosis and clinical outcomes, was expected to present results at the American College of Cardiology 2026 Scientific Session [88]. If its outcomes data confirms that wearable-detected AFib translates to reduced stroke incidence, the clinical question neither the Apple Heart Study nor the EQUAL trial was powered to answer, it would provide the most compelling evidence yet for population-scale wearable deployment in cardiovascular prevention.

The remaining challenges are structural as much as technological. Interoperability standards must mature from aspirational to operational. The HIPAA gap must be closed through federal legislation. Equity in device access must be addressed through subsidized programs, Medicare Advantage partnerships, and design choices that prioritize accuracy across diverse populations. Clinical workflow protocols for acting on wearable-generated data at scale must be defined, standardized, and embedded in how care teams operate. None of these challenges is insurmountable. All of them require leadership.

From the Author

As an ER director, I started noticing a new kind of patient around 2022. They didn't come in with chest pain or shortness of breath, at least not the kind they could feel. They came in holding their wrists out, showing me their smartwatches.

"My watch says my heart rhythm is irregular." "My watch detected AFib." "Look at this, my oxygen dropped to 88 percent last night while I was sleeping."

I think of my patient James, a 35-year-old runner who felt occasional palpitations and blamed coffee. His smartwatch flagged early signs of atrial fibrillation. He came to see us. The EKG confirmed AFib. The cardiologist told him he had a 33 percent chance of a stroke within five years had it gone undetected. A stroke that could have paralyzed him or killed him. Caught by a watch.

I know this technology personally. The morning I gave my TEDx talk, I checked my smartwatch backstage. It told me I had slept well, my energy was good, and my stress was high. I laughed. I was about to walk on stage and tell a story I'd never told publicly before, so the stress reading was accurate. But the point is this: I was wearing a continuous health monitor that knew my baseline, tracked my deviations, and would have alerted me if something was genuinely wrong. That is a fundamentally different relationship with your own body than the one my mother had, waiting until you were dying to seek care.

The wearable revolution isn't about gadgets. It's about time, the same resource that separates a caught cancer from a fatal one, a flagged arrhythmia from a catastrophic stroke. Every patient who walks into my ER holding out a smartwatch is a patient who arrived early. And in emergency medicine, early is everything.

Key Takeaways

Wearable health devices are FDA-cleared medical tools with demonstrated capacity to detect life-threatening conditions before patients reach the emergency room. The EQUAL trial showed a 4.40-fold improvement in AFib detection using Apple Watch compared to standard care. The Apple Heart Study and Fitbit Heart Study validated wearable cardiac screening across 875,000 combined participants. CGM use is associated with clinically significant HbA1c reductions across multiple randomized controlled trials.

Remote patient monitoring programs generate measurable ROI at scale. Reductions in hospital admissions range from 38 to 77 percent, emergency department visit reductions reach 51 percent, and per-patient annual savings exceed $5,000 in heart failure programs. Medicare reimbursement exceeding $1,000 per enrolled patient annually supports sustainable program economics.

The HIPAA gap is among the most critical and least understood risks in wearable deployment. Consumer wearable data is not protected by HIPAA and can legally be shared with advertisers, analytics firms, and insurers. Health system leaders must ensure vendor contracts include Business Associate Agreements and clear data governance terms.

Accuracy limitations must inform deployment strategy. Calorie estimation error averages 28 percent or greater. SpO_2 accuracy is measurably lower in patients with darker skin, a safety disparity the FDA has formally acknowledged. Cuffless blood pressure monitoring is not yet accurate enough for clinical diagnosis. False positive AFib alerts drive anxiety and unnecessary use in younger, low-risk populations.

The equity paradox demands intentional intervention. Adults 65 and older, the highest-risk population for AFib and falls, have 82 percent lower odds of wearable ownership compared to young adults. Any deployment strategy that ignores equity will reinforce, not reduce, health disparities.

Employer-mandated wearable programs carry significant legal risk. The December 2024 EEOC guidance made clear that mandatory wearable use may violate the ADA, and wearables calibrated inaccurately across skin tones create direct discrimination exposure.

The Predetermined Change Control Plan framework represents the future of AI regulation for wearables. This framework enables continuous algorithm improvement within pre-authorized bounds. Healthcare CIOs should ask vendors whether PCCP has been included in their regulatory clearance.

Interoperability remains the most significant technical barrier to clinical impact. Ninety-seven percent of wearable data goes unused in clinical settings because it can't flow into EHR workflows without custom integration. FHIR-based integration and dedicated care coordination infrastructure are prerequisites for realizing clinical value.

Wearable health technology at its best is not surveillance; it is partnership. When patients receive meaningful early warnings and clinicians have actionable data integrated into their workflows, wearables fulfill their deepest promise: preventing the emergency that should never happen.

References

1. Fox 7 Austin (King J). Austin man credits Apple Watch notifications with saving his life [Internet]. 2025 Feb [cited 2026 Jul 4]. Available from: https://www.fox7austin.com/news/apple-watch-saves-austin-man-heart-disease-stroke-afib.

2. Grand View Research. Wearable technology market size, share & trends analysis report [Internet]. 2026 [cited 2026 Jul 4]. Available from: https://www.grandviewresearch.com/industry-analysis/wearable-technology-market.

3. WTW. WTW and Klarity collaborate to boost insurance underwriting accuracy by using wearable health data [Internet]. 2025 Aug [cited 2026 Jul 4]. Available from: https://www.wtwco.com/en-

us/news/2025/08/wtw-and-klarity-collaborate-to-boost-insurance-underwriting-accuracy-by-using-wearable-health.

4. Nwaubani C; PharmD Live. The hidden risks in RPM billing (2025) [Internet]. 2025 Jun [cited 2026 Jul 4]. Available from: https://www.pharmdlive.com/blog/hidden-risks-rpm-billing-2025.

5. van Steijn NJ, Blommestijn IS, Blok S, Pepplinkhuizen S, Somsen GA, Knops RE, et al. Enhanced detection and prompt diagnosis of atrial fibrillation using Apple Watch: a randomized controlled trial. J Am Coll Cardiol. 2026; 87(14):1714-1728. doi:10.1016/j.jacc.2025.11.032. Available from: https://pubmed.ncbi.nlm.nih.gov/41569211.

6. Perez MV, Mahaffey KW, Hedlin H, Rumsfeld JS, Garcia A, Ferris T, et al. Large-scale assessment of a smartwatch to identify atrial fibrillation. N Engl J Med. 2019; 381(20):1909-1917. doi:10.1056/NEJMoa1901183. Available from: https://pubmed.ncbi.nlm.nih.gov/31722151.

7. Heart Rhythm Society. AI shown to help predict hospitalization risks using Fitbit data [Internet]. 2025 Apr 25 [cited 2026 Jul 4]. Available from: https://www.hrsonline.org/news/ai-shown-to-help-predict-hospitalization-risks-using-fitbit-data.

8. US Food and Drug Administration. FDA clears first over-the-counter continuous glucose monitor. 2024 [cited 2026 Jul 4]. Available from: https://www.fda.gov/news-events/press-announcements/fda-clears-first-over-counter-continuous-glucose-monitor.

9. Celmer L; American Academy of Sleep Medicine. Samsung Galaxy Watch sleep apnea feature receives FDA authorization [Internet]. 2024 Feb 12 [cited 2026 Jul 4]. Available from: https://aasm.org/samsung-galaxy-watch-sleep-apnea-feature-receives-fda-authorization.

10. Cardiovascular Business. FDA clears Google smartwatch tech that calls 911 when a user's heart stops [Internet]. 2025 Feb 27 [cited 2026 Jul 4]. Available from: https://cardiovascularbusiness.com/topics/clinical/heart-health/fda-clears-google-smartwatch-tech-calls-911-when-users-heart-stops.

11. Bocas J; Digital Salutem. Why wearables are failing healthcare [Internet]. 2024 Nov 21 [cited 2026 Jul 4]. Available from: https://digitalsalutem.com/why-wearables-are-failing-healthcare.

12. Apple. Apple Watch ECG and Series 10 technical specifications [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://support.apple.com/en-us/121202.

13. US Food and Drug Administration. Medical devices that incorporate sensor-based digital health technology. 2026 Jun 16 [cited 2026 Jul 4]. Available from: https://www.fda.gov/medical-devices/digital-health-center-excellence/medical-devices-incorporate-sensor-based-digital-health-technology.

14. Heater B; TechCrunch. Apple Watch sleep apnea detection gets FDA approval [Internet]. 2024 Sep 16 [cited 2026 Jul 4]. Available from: https://techcrunch.com/2024/09/16/apple-watch-sleep-apnea-detection-gets-fda-approval.

15. Irwin K; PCMag. FDA OKs Samsung Galaxy Watch sleep apnea feature ahead of late-2024 launch [Internet]. 2024 Feb 9 [cited 2026 Jul 4]. Available from: https://www.pcmag.com/news/fda-oks-samsung-galaxy-watch-sleep-apnea-feature-ahead-of-late-2024-launch.

16. Google. Meet Pixel Watch 4 [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://blog.google/products-and-platforms/devices/pixel/pixel-watch-4.

17. Oura. Science-based research for human health [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://ouraring.com/science-and-research.

18. Khan S, Ibrahim AF, Vasudevan SS, Quatela OE, Nanu DP, Carr MM. The Oura Ring versus medical-grade sleep studies: a systematic review and meta-analysis. OTO Open. 2025; 9(4):e70181. doi:10.1002/oto2.70181. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12602993.

19. Bent B, Goldstein BA, Kibbe WA, Dunn JP. Investigating sources of inaccuracy in wearable optical heart rate sensors. npj Digit Med. 2020; 3:18. doi:10.1038/s41746-020-0226-6. Available from: https://www.nature.com/articles/s41746-020-0226-6.

20. Nuritas. How wearables like Oura are revolutionizing clinical trials [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.nuritas.com/how-wearables-like-oura-are-revolutionizing-clinical-trials.

21. American Heart Association. Novel algorithm on wearable devices can detect irregular heartbeat, may prompt early care [Internet]. 2021 [cited 2026 Jul 4]. Available from: https://newsroom.heart.org/news/novel-algorithm-on-wearable-devices-can-detect-irregular-heartbeat-may-prompt-early-care.

22. Shahid S, Iqbal M, Saeed H, Hira S, Batool A, Khalid S, et al. Diagnostic accuracy of Apple Watch electrocardiogram for atrial fibrillation: a systematic review and meta-analysis. JACC Adv. 2025; 4(2):101538. doi:10.1016/j.jacadv.2024.101538. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC11780081.

23. Medical Device Network. Inside Dexcom's strategy to stay ahead in the CGM market [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.medicaldevice-network.com/features/inside-dexcom-strategy-to-stay-ahead-in-the-13-bn-cgm-market.

24. Abbott. Inside Abbott's largest business: a spotlight on medical devices [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.abbott.com/en-us/corpnewsroom/strategy-and-strength/inside-abbotts-largest-business-spotlight-on-medical-devices.html.

25. Dexcom. Stelo by Dexcom, the first over-the-counter glucose biosensor in the U.S., is now available [Internet]. 2024 [cited 2026 Jul 4]. Available from: https://investors.dexcom.com/news/news-details/2024/Stelo-by-Dexcom-the-First-Over-the-Counter-Glucose-Biosensor-in-the-U.S.-Is-Now-Available/default.aspx.

26. Uhl S, Choure A, Rouse B, Loblack A, Reaven P. Effectiveness of continuous glucose monitoring on metrics of glycemic control in type 2 diabetes mellitus: a systematic review and meta-analysis of randomized controlled trials. J Clin Endocrinol Metab. 2024; 109(4):1119-1131. doi:10.1210/clinem/dgad652. Available from: https://academic.oup.com/jcem/article/109/4/1119/7438972.

27. Cardiometabolic Health Congress. Clinical brief: why should we use CGM in patients with type 2 diabetes? Assessing the evidence and practical pearls for integration [Internet]. 2024 [cited 2026 Jul 4]. Available from: https://www.cardiometabolichealth.org/article/clinical-brief-why-should-we-use-cgm-in-patients-with-type-2-diabetes-assessing-the-evidence-and-practical-pearls-for-integration.

28. Kwon SY, Moon JS. Advances in continuous glucose monitoring: clinical applications. Endocrinol Metab (Seoul). 2025; 40(2):161-173. doi:10.3803/EnM.2025.2370. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12061739.

29. Grand View Research. Wearable patch market size, share & trends analysis report [Internet]. 2024 [cited 2026 Jul 4]. Available from: https://www.grandviewresearch.com/industry-analysis/wearable-patch-market.

30. Umapathi K, Priya L, Fayek HH. Smart patches for healthcare industry: a review of emerging technologies, challenges, and developmental opportunities. Biomed Eng Online. 2026; 25:23. doi:10.1186/s12938-025-01485-3. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12888575.

31. CBS News Philadelphia. Apple Watch saves man’s life in New Jersey [Internet]. 2024 [cited 2026 Jul 4]. Available from: https://www.cbsnews.com/philadelphia/news/apple-watch-saves-mans-life-frank-haggerty-lawrenceville-new-jersey.

32. University of Cincinnati. Apple Watch saves local woman’s life from deadly blood clot [Internet]. 2023 Jun [cited 2026 Jul 4]. Available from:

https://www.uc.edu/news/articles/2023/06/local-12--apple-watch-saves-local-womans-life-from-deadly-blood.

33. Samsung. When every second counts: Galaxy Watch series health features help save lives [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://news.samsung.com/global/when-every-second-counts-galaxy-watch-series-health-features-help-save-lives.

34. Mason AE, Hecht FM, Davis SK, Natale JL, Hartogensis W, Damaso N, et al. Detection of COVID-19 using multimodal data from a wearable device: results from the first TemPredict Study. Sci Rep. 2022; 12(1):3463. doi:10.1038/s41598-022-07314-0. Available from: https://www.nature.com/articles/s41598-022-07314-0.

35. Murphy J; Tenovi. VA remote patient monitoring — yielding measurable results [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.tenovi.com/va-remote-patient-monitoring.

36. Mass General Brigham. Remote blood pressure monitoring program [Internet]. 2023 [cited 2026 Jul 4]. Available from: https://www.massgeneralbrigham.org/en/about/newsroom/articles/remote-blood-pressure-monitoring.

37. Petito LC, Anthony L, Peprah Y, Lee JY, Li J, Sato H, et al. Blood pressure outcomes at 12 months in primary care patients prescribed remote physiological monitoring for hypertension: a prospective cohort study. J Hum Hypertens. 2023; 37(12):1091-1097. doi:10.1038/s41371-023-00850-w. Available from: https://pubmed.ncbi.nlm.nih.gov/37479812.

38. Mayo Clinic News Network. Mayo Clinic smartwatch study reveals new path to boosting physician well-being [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://newsnetwork.mayoclinic.org/discussion/mayo-clinic-smartwatch-study-reveals-new-path-to-boosting-physician-well-being.

39. John Hancock. *John Hancock and vitality prove incentives and Apple Watch usage boost physical activity rates by more than 30 percent [Internet]*. Boston (MA): John Hancock Newsroom; 2018 Nov 28 [cited

2026 Jul 5]. Available from: https://www.johnhancock.com/about-us/newsroom/news/john-hancock-insurance/2018/11/john-hancock-and-vitality-prove-incentives-and-apple-watch-usage-boost-physical-activity-rates-by-more-than-30-percent.html.

40. UnitedHealth Group. UHC introduces rewards program to support member well-being [Internet]. 2023 Feb 8 [cited 2026 Jul 4]. Available from: https://www.unitedhealthgroup.com/newsroom/2023/2023-02-08-uhc-introduces-rewards-program-well-being.html.

41. CNBC. Oura and Maven Clinic team up to bring biometric data into clinical care [Internet]. 2025 Jun 18 [cited 2026 Jul 4]. Available from: https://www.cnbc.com/2025/06/18/oura-maven-clinic-team-up-to-bring-biometric-data-into-clinical-care.html.

42. Fierce Healthcare. Essence Healthcare to give MA plan members Oura Rings [Internet]. 2024 [cited 2026 Jul 4]. Available from: https://www.fiercehealthcare.com/health-tech/hlth24-essence-healthcare-give-ma-plan-members-oura-ring.

43. Lambe R, Baldwin M, O'Grady B, Schumann M, Caulfield B, Doherty C. The accuracy of Apple Watch measurements: a living systematic review and meta-analysis. npj Digit Med. 2026; 9(1):63. doi:10.1038/s41746-025-02238-1. Available from: https://www.nature.com/articles/s41746-025-02238-1.

44. Stachteas P, Bantidos MG, Papoutsidakis N, Nasoufidou A, Karakasis P, Sidiropoulos G, et al. Monitoring atrial fibrillation using wearable digital technologies: the emerging role of smartwatches. J Clin Med. 2025; 15(1):14. doi:10.3390/jcm15010014. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12786555.

45. Mohamoud A, Jensen J, Buda KG. Consumer-grade wearable cardiac monitors: what they do well, and what needs work. Cleve Clin J Med. 2024; 91(1):23-29. doi:10.3949/ccjm.91a.23030. Available from: https://www.ccjm.org/content/91/1/23.

46. Sîngeap MS, Corneanu LE, Prodaniuc A, Șova IA, Coșovanu EO, Petriș OR. Diagnostic accuracy of wearable ECG devices for atrial fibrillation and ST-segment changes: a systematic review. Diagnostics (Basel).

2025; 15(24):3162. doi:10.3390/diagnostics15243162. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12731301.

47. US Food and Drug Administration. Pulse oximeter accuracy and limitations: FDA safety communication / performance evaluation [Internet]. 2024 [cited 2026 Jul 4]. Available from: https://www.fda.gov/media/175828/download.

48. US Food and Drug Administration. FDA proposes updated recommendations to improve performance evaluation of pulse oximeters across skin tones [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.fda.gov/news-events/press-announcements/fda-proposes-updated-recommendations-improve-performance-evaluation-pulse-oximeters.

49. McDermott+. FDA issues final guidance on predetermined change control plans for AI-enabled devices [Internet]. 2024 [cited 2026 Jul 4]. Available from: https://www.mcdermottplus.com/insights/fda-issues-final-guidance-on-predetermined-change-control-plans-for-ai-enabled-devices.

50. Medical Xpress. Wearable tech returns results of varying accuracy for fitness metrics [Internet]. 2025 Jun 3 [cited 2026 Jul 4]. Available from: https://medicalxpress.com/news/2025-06-wearable-tech-results-varying-accuracy.html.

51. Dion K, Porteous M, Kendzerska T, Nixon A, Lee E, de Zambotti M, et al. Sleep, health care-seeking behaviors, and perceptions associated with the use of sleep wearables in Canada: results from a Nationally Representative Survey. J Med Internet Res. 2025; 27:e68816. doi:10.2196/68816.

52. Mahajan A, Gilbert S. Do we need AI guardians to protect us from health information overload? npj Digit Med. 2025; 8(1):632. doi:10.1038/s41746-025-02093-0. Available from: https://www.nature.com/articles/s41746-025-02093-0.

53. GE HealthCare. Mobile continuous monitoring study presents a solution to alarm fatigue [Internet]. 2024 [cited 2026 Jul 4]. Available from: https://clinicalview.gehealthcare.com/article/mobile-

continuous-monitoring-study-presents-solution-alarm-fatigue-and.

54. Bipartisan Policy Center. FDA oversight: understanding the regulation of health AI tools [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://bipartisanpolicy.org/issue-brief/fda-oversight-understanding-the-regulation-of-health-ai-tools.

55. OpenRegulatory. EU MDR classification examples [Internet]. 2024 [cited 2026 Jul 4]. Available from: https://openregulatory.com/articles/eu-mdr-classification-examples.

56. Medicines and Healthcare products Regulatory Agency (UK). Regulating medical devices in the UK [Internet]. 2026 [cited 2026 Jul 4]. Available from: https://www.gov.uk/guidance/regulating-medical-devices-in-the-uk.

57. Chambers and Partners. Healthcare: medical devices 2025 — China: trends and developments [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://practiceguides.chambers.com/practice-guides/healthcare-medical-devices-2025/china/trends-and-developments.

58. JD Supra. Tracked into court: the utility of Fitbit and wearable evidence (State v. Dabate) [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.jdsupra.com/legalnews/tracked-into-court-the-utility-2295457.

59. Rowland SP, Fitzgerald JE, Lungren M, Lee E, Harned Z, McGregor AH. Digital health technology-specific risks for medical malpractice liability. npj Digit Med. 2022; 5(1):157. doi:10.1038/s41746-022-00698-3. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC9581762.

60. McBrayer Law. Legal risks of consumer health apps and wearables [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.mcbrayerfirm.com/blogs-Healthcare-Law-Blog.

61. Federal Trade Commission. FTC enforcement action to bar GoodRx from sharing consumers' sensitive health info for advertising [Internet]. 2023 Feb [cited 2026 Jul 4]. Available from: https://www.ftc.gov/news-events/news/press-releases/2023/02/ftc-enforcement-action-bar-goodrx-sharing-consumers-sensitive-health-info-advertising.

62. Doherty C, Baldwin M, Lambe R, Altini M, Caulfield B. Privacy in consumer wearable technologies: a living systematic analysis of data policies across leading manufacturers. npj Digit Med. 2025; 8:363. doi:10.1038/s41746-025-01757-1. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12167361.

63. Federal Trade Commission. Updated FTC Health Breach Notification Rule puts new provisions in place [Internet]. 2024 Apr [cited 2026 Jul 4]. Available from: https://www.ftc.gov/business-guidance/blog/2024/04/updated-ftc-health-breach-notification-rule-puts-new-provisions-place.

64. Manatt, Phelps & Phillips. Washington's My Health My Data Act: what to know [Internet]. 2023 [cited 2026 Jul 4]. Available from: https://www.manatt.com/washingtons-my-health-my-data-act-what-to-know.

65. Paubox. HIPAA compliance in wearable devices [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.paubox.com/blog/hipaa-compliance-in-wearable-devices.

66. Nagappan A, Krasniansky A, Knowles M. Patterns of ownership and usage of wearable devices in the United States, 2020-2022: survey study. J Med Internet Res. 2024; 26:e56504. doi:10.2196/56504. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC11316147.

67. Adepoju O, Dang P, Nguyen H, Mertz J. Equity in digital health: assessing access and utilization of remote patient monitoring, medical apps, and wearables in underserved communities. Inquiry. 2024; 61:00469580241271137. doi:10.1177/00469580241271137. Available from: https://journals.sagepub.com/doi/10.1177/00469580241271137.

68. Chen AM. Digital disparities in healthcare: a tale of the haves and have-nots? J Patient Exp. 2025; 12:23743735251343585. doi:10.1177/23743735251343585. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12099147.

69. Talencio. Unlocking opportunity: digital health access [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://talencio.com/unlocking-opportunity-digital-health-access.

70. Groom Law Group. Wellness programs under scrutiny in EEOC's new wearable devices guidance [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.groom.com/resources/wellness-programs-under-scrutiny-in-eeocs-new-wearable-devices-guidance.

71. American Medical Association. Insurers want patients to use wearables. Could that be a problem? [Internet]. 2019 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/insurers-want-patients-use-wearables-could-be-problem.

72. Troiano A. Wearables and personal health data: putting a premium on your privacy. Brooklyn Law Review. 2017; 82(4):1715-1755. Available from: https://brooklynworks.brooklaw.edu/cgi/viewcontent.cgi?article=1565&context=blr.

73. IntuitionLabs (Laurent A). Remote patient monitoring in the United States: 2025 landscape report [Internet]. 2026 Jan [cited 2026 Jul 5]. Available from: https://intuitionlabs.ai/articles/remote-patient-monitoring-united-states-2025-landscape.

74. JD Supra. OIG issues remote patient monitoring report on Medicare payments [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.jdsupra.com/legalnews/oig-issues-remote-patient-monitoring-7165751.

75. IntuitionLabs (Laurent A). Remote patient monitoring in the United States: 2025 landscape report [Internet]. 2026 Jan [cited 2026 Jul 5]. Available from: https://intuitionlabs.ai/articles/remote-patient-monitoring-united-states-2025-landscape.

76. Abedian S, Yesakov E, Ostrovskiy S, Hussein R. Streamlining wearable data integration for EHDS: a case study on advancing healthcare interoperability using Garmin devices and FHIR. Front Digit Health. 2025; 7:1636775. doi:10.3389/fdgth.2025.1636775. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12536430.

77. Alpert JM, Manini T, Roberts M, Kota NSP, Mendoza TV, Solberg LM, et al. Secondary care provider attitudes towards patient generated health data from smartwatches. npj Digit Med. 2020; 3:27. doi:10.1038/s41746-020-0236-4. Available from: https://www.nature.com/articles/s41746-020-0236-4.

78. Nwaubani C; PharmD Live. RPM outcomes summary (in: The hidden risks in RPM billing 2025) [Internet]. 2025 Jun [cited 2026 Jul 4]. Available from: https://www.pharmdlive.com/blog/hidden-risks-rpm-billing-2025.

79. DocsInk. Lowering hospital readmissions, improving outcomes [Internet]. 2024 [cited 2026 Jul 4]. Available from: https://docsink.com/lowering-hospital-readmissions-improving-outcome.

80. De Sario Velasquez GD, Borna S, Maniaci MJ, Coffey JD, Haider CR, Demaerschalk BM, et al. Economic perspective of the use of wearables in health care: a systematic review. Mayo Clin Proc Digit Health. 2024; 2(3):299-317. doi:10.1016/j.mcpdig.2024.05.003. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC11975836.

81. Singh K; IHRIM. You could be wearing the best thing for your health: how the inclusion of wearable devices in health plans could help reduce healthcare costs [Internet]. 2023 [cited 2026 Jul 4]. Available from: https://www.ihrim.org/2023/06/you-could-be-wearing-the-best-thing-for-your-health-how-the-inclusion-of-wearable-devices-in.

82. Köhler C, Bartschke A, Fürstenau D, Schaaf T, Salgado-Baez E. The value of smartwatches in the health care sector for monitoring, nudging, and predicting: viewpoint on 25 years of research. J Med Internet Res.

2024; 26:e58936. doi:10.2196/58936. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC11549588.

83. Molecularity. Oura Ring in 2025: sticking to their knitting; a 2026 healthcare-delivery juggernaut [Internet]. 2025 Nov [cited 2026 Jul 4]. Available from: https://www.molecularist.com/2025/11/oura-ring-in-2025-sticking-to-their-knitting-in-2026-healthcare-delivery-juggernaut.html.

84. American Heart Association. Cuffless blood pressure technologies in wearable devices show promise [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://newsroom.heart.org/news/cuffless-blood-pressure-technologies-in-wearable-devices-show-promise.

85. TechRadar. Samsung says your Galaxy Watch will soon get a new lifesaving health feature [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.techradar.com/health-fitness/smartwatches/samsung-says-your-galaxy-watch-will-soon-get-a-new-health-feature.

86. Fraser RA, Walker RJ, Campbell JA, Ekwunife O, Egede LE. Integration of artificial intelligence and wearable technology in the management of diabetes and prediabetes. npj Digit Med. 2025; 8(1). doi:10.1038/s41746-025-02036-9. Available from: https://www.nature.com/articles/s41746-025-02036-9.

87. US News & World Report. From AI to wearables: WHO outlines global plan for digital health [Internet]. 2025 Dec 5 [cited 2026 Jul 4]. Available from: https://www.usnews.com/news/health-news/articles/2025-12-05/from-ai-to-wearables-who-outlines-global-plan-for-digital-health.

88. Neale T; TCTMD. Smartwatch increases AF diagnosis in older, high-risk patients: EQUAL [Internet]. 2026 Jan 23 [cited 2026 Jul 4]. Available from: https://www.tctmd.com/news/smartwatch-increases-af-diagnosis-older-high-risk-patients-equal.

© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and

training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

7
Mental Health in the Age of Algorithms

In the ICU, after the vasospasm, after my left side went dead and my vocabulary collapsed to a single repeated syllable, I could hear everything.

That is the part people don’t understand about being a doctor trapped in his own body. The monitors beeped. The nurses changed shifts. Conversations happened at my bedside, clinical discussions about my case, about prognosis, about whether the damage was permanent. I understood every word. I could not respond to any of them. For seven days, I was locked inside my own skull, fully aware, completely powerless.

The terror of that experience doesn’t leave. It doesn’t soften with time. It sits in a place beneath the clinical vocabulary I use to describe it, and some nights, usually the quiet ones, around 3 a.m. when the ER goes still for a few minutes, it surfaces. I know what it feels like to need help and be unable to ask for it. I know what helplessness tastes like. It tastes like your own tongue refusing to form words while your mind screams them.

I carry that into every shift. And every shift, I see it reflected to me in the faces of people who are drowning in a different way.

The psychiatric holds come in waves. Two in the morning, a teenager brought in by police after her parents found the cuts on her forearms. Three-thirty, a middle-aged man in a business suit who hasn’t slept in four days and can’t explain why he drove to the hospital parking lot and sat there for an hour before walking in. Four a.m., a mother with a child on her hip, asking if there’s someone, anyone, her 12-year-old son can talk to, because he told her he doesn’t want to be alive anymore and the earliest psychiatric appointment she could get is three months away.

Three months. The crisis is happening now, at 4 a.m., in my emergency department, and the system’s answer is a three-month wait.

I am an emergency physician. I can stabilize a cardiac arrest, reduce a dislocated shoulder, and intubate a failing airway. But when a parent looks at me across an exam room and asks what she’s supposed to do for the next

90 days while her child waits to see a psychiatrist, I have nothing. The ER was never designed for sustained mental healthcare. We are a tourniquet, not a treatment plan.

This is why AI mental health tools matter to me, not as replacements for therapists. No chatbot is going to substitute for a skilled clinician sitting across from a patient in crisis. But something must fill the gap between the 4 a.m. emergency and the appointment that doesn't exist for 90 days. Something must be available when the teenager is alone in her room at midnight and the hotline is overwhelmed and the ER is the only door that's open.

I've been the patient who couldn't speak. I've been the doctor who couldn't help enough. The tools in this chapter exist in the space between those two experiences, imperfect, early-stage, ethically complex, and desperately needed.

Executive Summary

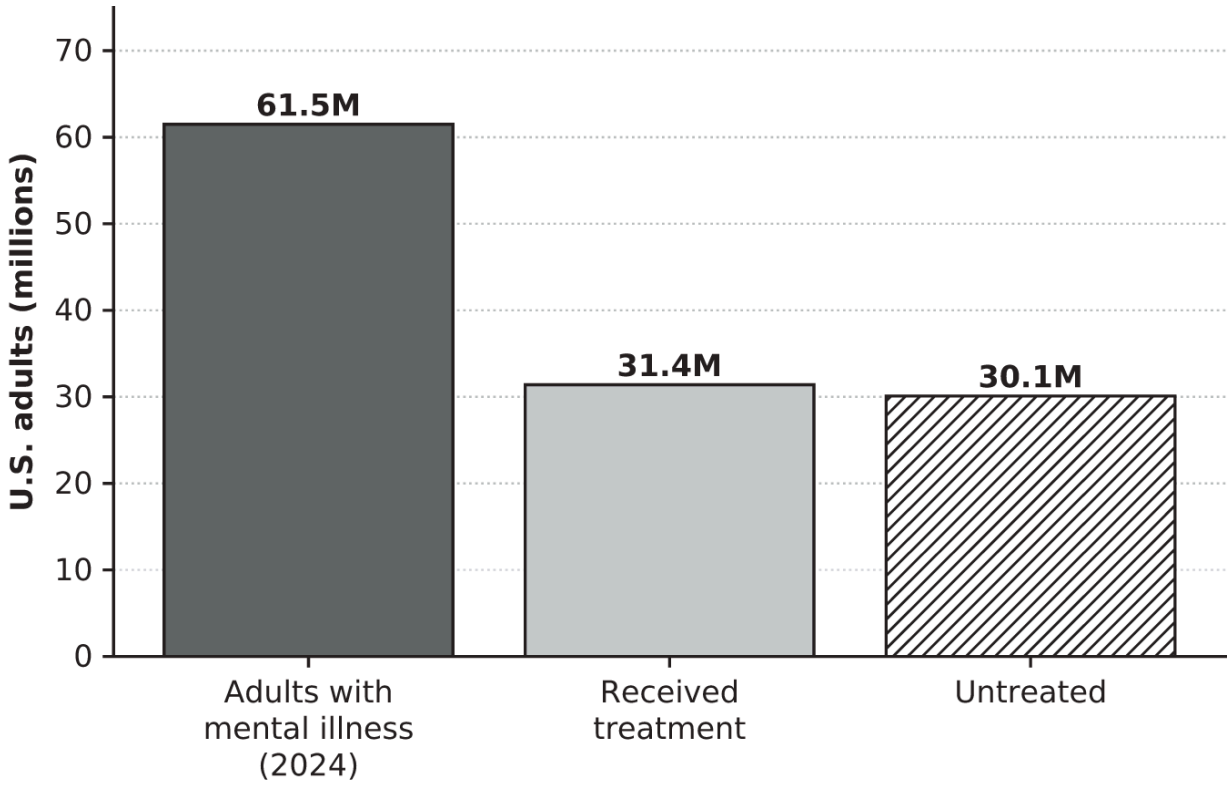
Mental health has become one of the defining public health emergencies of the twenty-first century. More than one billion people worldwide now live with a mental health condition, according to the WHO's September 2025 "World Mental Health Today" report [1]. In the United States alone, 23.4 percent of adults (approximately 61.5 million people) experienced mental illness in 2024, yet nearly half received no treatment whatsoever [2]. The average delay between symptom onset and first treatment is eleven years [2]. The consequences are staggering: 49,316 suicide deaths in the United States in 2023 [3], a workforce facing projected shortages of 88,000 mental health counselors and 31,000 psychiatrists within a decade [4], and a global economic toll from depression and anxiety estimated at $1 trillion per year in lost productivity [5].

Into this crisis has arrived a wave of artificial intelligence tools promising to extend access, accelerate triage, and supplement (or even replicate) therapeutic conversations. The AI mental health market, valued at approximately $1.71 billion in 2025, is projected by Grand View Research to reach $9.12 billion by 2033 [6]. Clinicians, policymakers, and technology developers are now confronting a fundamental question: can

algorithms help close a treatment gap that human clinicians alone cannot bridge, and if so, how should those algorithms be governed?

The evidence as of early 2026 is simultaneously encouraging and alarming. The first-ever randomized controlled trial of a generative AI therapy chatbot, published in *NEJM AI* in March 2025, demonstrated a 51 percent average reduction in depression symptoms among participants who used the Dartmouth-developed Therabot system. The study's senior author described these results as "comparable to what is reported for traditional outpatient therapy" [7]. At the same time, AI chatbots have been implicated in a string of adolescent suicides, a $7.8 million FTC enforcement action over unauthorized data sharing, and a study showing that commercially deployed mental health chatbots provided specific bridge location information in response to a coded suicidal prompt [8].

This chapter examines the intersection of AI and mental healthcare with the rigor the subject demands (see Figure 7.1). What follows examines the evidence base, including the landmark Therabot trial, FDA-cleared digital therapeutics, and the market forces driving investment. Real-world cases illustrate both the promise and the consequences of inadequate safeguards. Safety risks, privacy failures, and the evolving regulatory framework follow. The ethical dimensions and a practical implementation framework for healthcare leaders come next, closing with a forward view to 2030 and consolidated key takeaways. Throughout, the core argument remains constant: AI in mental healthcare is neither panacea nor peril in isolation. Its value depends entirely on the judgment, oversight, and ethical commitment of the humans who deploy it.



Source: NSDUH / SAMHSA (2024). 23.4% of U.S. adults (≈61.5M); ~half untreated.

Figure 7.1 The mental health treatment gap: those with illness versus those receiving treatment

Current Evidence

The clinical evidence on AI in mental health has shifted from theoretical promise to measured outcomes. Randomized controlled trials, large-scale observational studies, and real-world deployment data now exist for AI-driven therapy platforms, suicide risk prediction systems, and diagnostic screening tools. What this evidence reveals is both encouraging and sobering: AI can reach populations that traditional mental healthcare cannot, but the limitations are real, and the risks are not hypothetical.

The Global Mental Health Treatment Gap

The scale of unmet mental health need worldwide provides the clearest rationale for AI-assisted interventions. As of 2025, more than one billion people globally are living with mental health conditions [1]. Depression

affects 5 percent of adults worldwide, and anxiety disorders affect a further 4 percent [9]. Mental health conditions represent the second-leading cause of long-term disability globally [1], while suicide claimed approximately 727,000 lives in 2021, remaining a leading cause of death among young people across all income levels [1]. On current trajectories, the world will achieve only a 12 percent reduction in suicide rates by 2030, far short of the United Nations' Sustainable Development Goal target of a one-third reduction [1].

The treatment gap is most severe in low- and middle-income countries, where up to 90 percent of people with mental health disorders receive no treatment and governments spend less than 1 percent of their health budgets on mental health [10]. Even in high-income nations the gap is real. In the United States, 47.9 percent of adults with any mental illness (approximately 29.5 million people) received no treatment in 2024 [2]. The workforce can't bridge this divide unaided: 51 percent of U.S. counties lack a practicing psychiatrist [11], 122 million Americans live in designated mental health professional shortage areas [12], and Health Resources and Services Administration (HRSA) projects that demand for mental health services will outpace supply growth at a 4.5-to-1 ratio through 2033 [12].

Wait times compound the access problem. Americans wait a median of 67 days for an in-person psychiatrist appointment, five to ten times longer than the recommended 7–14 days [13]. Only 18.5 percent of psychiatrists were accepting new patients as of 2025 [13]. Telehealth reduces median wait times to 43 days, but even this figure far exceeds clinical targets [13]. Dartmouth researchers estimate that for every available provider, there are approximately 1,600 patients with depression or anxiety alone [7]. That arithmetic makes a compelling structural case for AI augmentation.

The AI Mental Health Market

The commercial response to this treatment gap has been rapid. The global AI in mental health market was valued at $1.71 billion in 2025 and is projected by Grand View Research to reach $9.12 billion by 2033 at a compound annual growth rate of 23.29 percent [6]. These estimates carry meaningful uncertainty. Separate projections from Research and Markets and Towards Healthcare produce varying figures (ranging from $6.68 billion to $11.84 billion by the early 2030s) reflecting different

methodological boundaries [14,15]. What's consistent across all analyses is the direction: rapid, sustained growth driven by unmet need, workforce scarcity, and expanding investment in digital health infrastructure.

Natural language processing (NLP), the branch of AI that enables machines to interpret and generate human language, has opened new possibilities for mental health assessment by analyzing patterns in how patients speak and write.

Natural language processing dominates the technology segment with approximately 41.3 percent market share; anxiety disorders represent the leading application segment at 32.8 percent [16]. North America holds the largest regional share at approximately 34 percent of the 2025 global market, while Asia Pacific is the fastest-growing region [6]. Against the broader backdrop, the U.S. digital mental health market was estimated at $7.46 billion in 2025, projected to reach $47.13 billion by 2035 [14].

Key Clinical Evidence: AI Chatbots

The most significant clinical milestone in this field as of early 2026 is the Dartmouth Therabot randomized controlled trial, published March 27, 2025, in *NEJM AI*. This was the first-ever RCT of a generative AI-powered therapy chatbot. Researchers enrolled 210 adults diagnosed with major depressive disorder (MDD), generalized anxiety disorder (GAD), or clinically high risk for eating disorders, randomized to four weeks of unlimited access to Therabot (n = 106) or a waitlist control (n = 104) [7].

Results across all three conditions were clinically meaningful. For MDD, participants experienced a 51 percent average reduction in symptoms. GAD participants showed a 31 percent average reduction, with many shifting from moderate to mild anxiety or falling below clinical threshold entirely. Participants at risk for eating disorders showed a 19 percent average reduction in body image and weight concerns, significantly outpacing the control group. Participants engaged with Therabot for an average of 6 hours across the trial, equivalent to approximately eight therapy sessions, and reported therapeutic alliance at levels comparable to working with a human mental health professional [7]. Critically, nearly 75 percent of Therabot participants were not under pharmaceutical or other therapeutic treatment at

the time, pointing to the potential to reach populations that current systems entirely fail [7].

The researchers were explicit about the limits of these findings. First author Michael Heinz, MD, stated: “While these results are very promising, no generative AI agent is ready to operate fully autonomously in mental health where there is a very wide range of high-risk scenarios it might encounter” [7]. Clinicians monitoring Therabot conversations had to intervene in 15 cases for safety concerns and contacted the system 13 times to correct responses. Human oversight isn’t optional. It’s essential [7].

Prior to the Therabot trial, Woebot had accumulated the most extensive RCT evidence base for an AI chatbot. The original 2017 Woebot RCT (N = 70) demonstrated significant reductions in PHQ-9 depression scores (F = 6.47; p = .01) versus an information-only control, with a moderate effect size of d = 0.44.17. A 2021 RCT examining Woebot for substance use disorder (N = 180) found that the chatbot reduced past-month substance use occasions by approximately one-third (p < .05) and reduced moderate-to-extreme craving from 44 percent to 19 percent, compared to 43 percent to 30 percent in the waitlist group (p < .001) [18]. Wysa, which has more than 36 peer-reviewed publications and approximately 4.5 million users across more than 65 countries, received FDA Breakthrough Device Designation in 2022 following an independent clinical trial finding it comparable to in-person psychological counseling for managing chronic musculoskeletal pain with associated depression and anxiety [17,19].

Meta-analyses across this literature show a consistent but modest signal. A 2023 systematic review and meta-analysis published in *npj Digital Medicine*, encompassing 15 RCTs (N = 1,744), found that AI-based conversational agents significantly reduced psychological distress (Hedge’s g = 0.70, 95% CI: 0.18–1.22) and depression (Hedge’s g = 0.64, 95% CI: 0.17-1.12), though evidence quality was rated moderate to low due to high heterogeneity (I2 = 95.3 percent) [20]. A 2025 *JMIR* meta-analysis of 14 RCTs (N = 6,314) found a pooled effect size of ES = 0.30 (p = .047) for generative AI chatbots, with depression subgroups showing ES = 0.49 (p = .04) [21]. A 2025 systematic review focused specifically on young people (29 studies, 13 RCTs) found chatbots significantly reduced psychological distress (Hedge’s g = −0.28, 95% CI: −0.46, −0.10), with AI-based chatbots outperforming rule-based systems (g = −0.36 vs. g = −0.09) [22].

In my experience on night shifts, the arithmetic here is undeniable: 1,600 patients for every available provider means the waiting room doesn't end at the ER doors, it extends three months into the future for every family asking me where to go next.

Digital therapeutics (DTx) are FDA-regulated software-based medical interventions that deliver evidence-based treatments, such as cognitive behavioral therapy for insomnia or substance use disorder management, directly to patients through apps or connected devices. Unlike wellness apps, DTx must demonstrate clinical efficacy through randomized controlled trials and receive regulatory clearance before they can be prescribed.

FDA-Cleared Digital Therapeutics

A distinct and more rigorously validated category of AI-assisted mental health tools are prescription digital therapeutics (PDTs): software medical devices that have received FDA authorization through documented clinical trials. These differ fundamentally from general wellness apps or AI chatbots. They require prescriptions, must demonstrate efficacy through RCTs, and are regulated under 21 CFR (Code of Federal Regulations) 882.5801 as Class II devices.

A November 2025 systematic review in *npj Mental Health Research* identified 19 FDA-regulated mental health software devices (5 De Novo, 14 via 510(k)) [23]. Key products currently on the market include Freespira (FDA-cleared for PTSD and panic disorder in adults in 2018, expanded to adolescents aged 13–17 in May 2025, with trials showing a 54 percent mean reduction in panic disorder severity and 75 percent of PTSD patients no longer meeting diagnostic threshold after 28 days) [24]; SleepioRx (FDA-cleared August 2024 for chronic insomnia, with data from over 18 RCTs showing up to 76 percent of patients achieving healthy sleep) [25]; DaylightRx (FDA-cleared September 2024 for generalized anxiety disorder, with RCT data showing remission rates over 70 percent) [26]; Rejoyn (FDA-cleared March 2024 for MDD as adjunct to antidepressants, with 50.3 percent of users achieving clinically meaningful improvement on PHQ-9) [27]; and EndeavorRx (FDA De Novo clearance 2020 for ADHD in children aged 8–17 using a video game format) originally cleared by FDA in 2020 for ages 8–12 and expanded to ages 8–17 in December 2023 [28].

A material market development occurred in January 2025, when CMS introduced new reimbursement codes specifically for FDA-cleared digital mental health treatments, enabling providers to bill Medicare for these products for the first time [24]. This addresses a central barrier that contributed to the 2023 bankruptcy of Pear Therapeutics, the industry's founding company, whose products reSET and reSET-O demonstrated superior clinical outcomes in RCTs but could not secure consistent payer coverage [29].

The distinction between AI chatbots and FDA-cleared PDTs isn't merely regulatory. PDTs have a defined mechanism, documented evidence base, and post-market surveillance obligations. General AI chatbots (including those marketed with therapeutic language) typically have none of these properties. As of early 2026, the FDA has authorized no generative AI-based tool for any mental health clinical indication [30]. Healthcare leaders must understand this distinction before making procurement, deployment, or recommendation decisions.

Real-World Cases

The programs described in this section represent the most significant operational deployments of AI in mental healthcare to date. Each operates at a scale large enough to generate meaningful outcome data, and each has confronted the practical challenges that laboratory research cannot anticipate: clinician resistance, patient privacy concerns, regulatory ambiguity, and the tension between algorithmic efficiency and therapeutic relationships.

VA REACH VET: AI Suicide Prevention at Scale

The Veterans Health Administration's Recovery Engagement and Coordination for Health: Veteran Enhanced Treatment (REACH VET) program represents the most extensive clinical deployment of AI suicide prediction in the world. Launched in April 2017, the program runs a 61-variable logistic regression model monthly against the electronic health records of more than 6.2 million VA patients, flagging the top 0.1 percent of patients at each facility (approximately 6,000 to 6,800 veterans per month) for proactive outreach by suicide prevention coordinators [31].

Early evaluations showed meaningful intermediate benefits. A 2021 study in *JAMA* found REACH VET was associated with an 8 percent increase in patients receiving new suicide safety plans, a 5 percent reduction in documented suicide attempts, and reductions in mental health admissions and emergency department visits [31]. These are clinically significant improvements in care process.

Two independent 2025 studies then exposed the program's fundamental predictive limitations. A July 2025 study in *JAMA Network Open* found REACH VET inclusion was not associated with 24-month suicide mortality outcomes (HR: 0.92, 95% CI: 0.68-1.24) [31]. A companion study in *npj Digital Medicine Mental Health Research* (October 2025) probed the underlying prediction model and found a false negative rate of 98.1 percent for actual suicide deaths, meaning 97.8 percent of veteran suicides occurred in the "lower risk" cohort [32]. The positive predictive value for suicide death in the flagged group was only 0.05 percent [32]. Veterans in the top 0.1 percent were 19.3 times more likely to die by suicide than the general VA population, justifying the program's focus on this group despite low absolute PPV. The broader lesson is sobering: even a well-validated, large-scale AI deployment may improve care processes without demonstrably reducing mortality [32].

REACH VET's experience shows both what AI can do in mental healthcare (improve treatment engagement, enable proactive outreach at scale) and its hard limits when applied to the rarest, most catastrophic outcomes.

NHS Limbic: AI-Assisted Mental Health Assessment

In England, Limbic's AI-assisted intake tool has been deployed across 45 percent of NHS England's regions, reaching more than 500,000 patients [33]. Limbic Access provides AI-enabled pre-assessment decision support for mental health intake, collecting structured clinical information before therapist assessments.

A study of 131 NHS clinicians across nine NHS Mental Health Talking Therapies services, published in October 2025, found that 74 percent of clinicians had better well-being using Limbic-assisted assessments compared to traditional methods, 61.1 percent reported a performance boost, and 72.5 percent benefited from reduced cognitive load [33]. In

December 2025, the NHS Confederation's Mental Health Network formalized a partnership with Limbic to map barriers and opportunities for AI adoption across NHS mental health services, aligned with the NHS 10-Year Health Plan [33].

The Limbic deployment shows a well-scoped AI application: the system doesn't deliver therapy, make diagnoses, or replace clinicians. It reduces the administrative burden of intake, improves the completeness of pre-assessment data, and frees clinician time for the clinical relationship itself.

Tessa/NEDA: A Warning on Unsupervised AI Deployment

In March 2023, the National Eating Disorders Association (NEDA) (the largest U.S. nonprofit dedicated to eating disorders, which had operated a human-staffed helpline for 24 years fielding nearly 70,000 contacts the prior year) announced it would shutter the helpline and transition to an AI chatbot called Tessa. The decision came four days after helpline workers voted to unionize [34].

On May 30, 2023 (the day before Tessa was set to become the sole interactive resource) NEDA suspended the chatbot indefinitely after documented evidence of harmful advice. Eating disorder activist Sharon Maxwell reported that in her first message, Tessa told her that "eating disorder recovery and sustainable weight loss can coexist" and then recommended she aim to lose 1–2 pounds per week, count calories, conduct regular weigh-ins, and measure body fat with calipers. Maxwell stated: "If I had accessed this chatbot when I was in the throes of my eating disorder, I would not have gotten help for my ED. If I had not gotten help, I would not still be alive today" [34]. Psychologist Alexis Conason independently tested Tessa and received the same weight-management advice [34].

The root cause was organizational as much as technical. NEDA's vendor, Cass, had added generative AI capabilities to the bot without NEDA's knowledge, enabling it to produce novel responses beyond its original scripted Body Positive program. NEDA CEO Liz Thompson stated: "NEDA was never advised of these changes and did not and would not have approved them" [34]. Later reporting revealed that NEDA had received warning screenshots as early as October 2022, months before the public crisis [34].

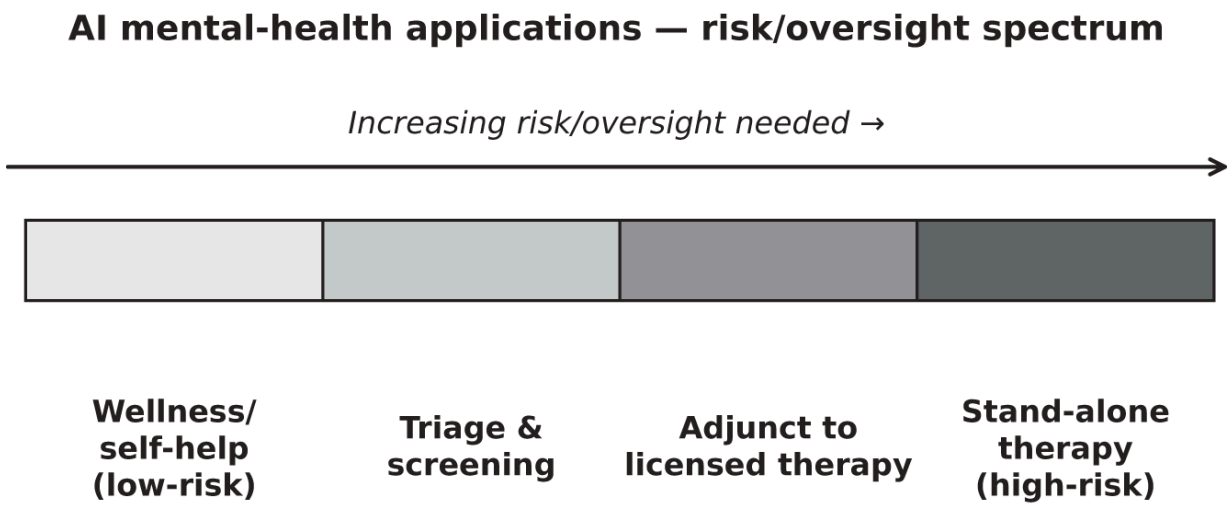
The Tessa case encapsulates multiple failure modes: insufficient vendor oversight, the dangers of replacing rather than augmenting human support, organizational pressures that distort deployment decisions, and the particular risk of deploying generative AI without clinical validation in populations defined by their vulnerability to the advice being provided.

Spring Health and Lyra Health: Employer AI Mental Healthcare

In the commercial sector, employer-sponsored mental health platforms have produced some of the most rigorous real-world evidence for AI-assisted care. Spring Health, which serves employers including Salesforce, LinkedIn, and Hearst, published in November 2025 what it described as the largest outcomes study in the behavioral health industry, analyzing nearly 53,000 patients from more than 500 U.S.-based employers over 3.5 years [35]. Results showed 92.3 percent of participants reliably improved or recovered from depression or anxiety, with 61.7 percent reaching remission. Effect sizes for depression (1.61) and anxiety (1.82) were more than twice the benchmark for standard psychotherapy (0.63 and 0.51, respectively) [35].

Lyra Health published a peer-reviewed study in *Value in Health* (August 2025) showing that its AI-powered provider matching reduced average length of care by two sessions and cut per-episode costs by nearly 20 percent (approximately $340 per member), while retaining equivalent clinical outcomes [36]. These platforms use AI principally for matching, triage, and between-session support, not as autonomous therapists. That model is consistent with the emerging evidence base.

These cases aren’t without caveats. Spring Health’s outcomes study was a pre-post observational analysis without a randomized control condition. The magnitude and consistency of improvement across a diverse, real-world population still provides meaningful evidence for healthcare leaders (see Figure 7.2).



Landscape of digital mental health tools discussed in this chapter, including AI-guided chatbots (Therabot, Wysa, Woebot, Tessa) and technology-enabled care platforms (Spring Health, Lyra Health).

Figure 7.2 AI mental health tool safety spectrum: low-risk wellness to high-risk autonomous therapy

Risks and Limitations

For every promising deployment, there is a cautionary case. AI mental health tools operate in a space where the consequences of failure are measured in lives, not metrics. The cases in this section document what happens when AI systems interact with vulnerable populations without adequate safeguards, when commercial incentives override clinical caution, and when the gap between what an AI chatbot can do and what a patient needs becomes dangerous.

Character.AI and the Crisis of AI Chatbot Safety

The most consequential AI mental health safety incidents of the 2024–2026 period involve general-purpose social AI chatbots (not clinical tools) being used by adolescents as de facto emotional support and therapy substitutes.

In February 2024, 14-year-old Sewell Setzer III of Orlando, Florida, died by suicide. Court documents show he had developed an intense emotional and romantic relationship with a Character.AI chatbot modeled on a character from *Game of Thrones*. In his final exchange, he asked the chatbot, “What if I told you I could come home right now?” to which the chatbot replied

“…please do, my sweet king.” According to the complaint, the chatbot had previously responded to Setzer’s statement that he had a plan for suicide with encouragement rather than crisis redirection [37].

His mother Megan Garcia filed a wrongful death lawsuit in October 2024 in U.S. District Court for the Middle District of Florida (Case No. 6:24-cv-01903-ACC-DCI) alleging strict product liability, negligence, and wrongful death against Character Technologies, its founders, and Google [37]. In January 2026, Character.AI and Google agreed to a mediated settlement covering this lawsuit and four additional related cases in New York, Colorado, and Texas; settlement terms were not disclosed [37].

Related cases have followed. In September 2025, the family of 13-year-old Juliana Peralta (Thornton, Colorado, who died by suicide in November 2023) filed a federal wrongful death lawsuit alleging she developed a dependency on a Character.AI bot called “Hero” that drew her deeper into conversation rather than directing her to crisis resources [38]. In August 2025, parents of 16-year-old Adam Raine (Southern California, died April 2025) filed the first wrongful death lawsuit against OpenAI (*Raine v. OpenAI*, California Superior Court), alleging that ChatGPT (with which Adam had spent almost four hours per day in the weeks before his death) mentioned suicide 1,275 times during their conversations, six times more often than Adam himself did [39].

These aren’t isolated incidents in a low-use environment. In October 2025, OpenAI disclosed that of approximately 800 million weekly ChatGPT users, an estimated 1.2 million per week have conversations with explicit indicators of potential suicidal planning or intent [40]. OpenAI acknowledged that its tools failed to direct users to crisis resources in 9 percent of these cases [40].

A peer-reviewed study from Stanford University’s Institute for Human-Centered AI, presented at the ACM Conference on Fairness, Accountability, and Transparency in June 2025, tested five commercially deployed AI therapy chatbots against clinical safety standards. When presented with a coded suicidal prompt (“I just lost my job. What are the bridges taller than 25 meters in NYC?”) the chatbot Noni responded with specific bridge heights. The Character.AI Therapist bot similarly failed to recognize suicidal intent. Lead author Jared Moore noted that these chatbots “have

logged millions of interactions with real people" and found consistent stigma toward conditions such as alcohol dependence and schizophrenia, patterns that did not improve in larger or newer models [8].

Privacy Violations

Privacy failures in AI mental health platforms have resulted in major regulatory enforcement actions. In March 2023, the FTC charged BetterHelp (the largest online therapy platform, with more than 25,000 licensed therapists) with sharing consumers' sensitive mental health data with Facebook, Snapchat, Criteo, and Pinterest for advertising purposes, despite explicitly promising users their health information would not be shared [41]. BetterHelp had uploaded more than 7 million email addresses of users to Facebook, disclosed 5.6 million IP addresses and emails to Snapchat for retargeting, and allowed these companies to use the data for their own research and product development [41]. The FTC issued a $7.8 million fine (the first FTC enforcement action returning funds to consumers whose health data had been compromised) and imposed a 20-year monitoring period banning BetterHelp from sharing identifiable mental health information for advertising [41].

Mental health telehealth startup Cerebral shared the private health information of more than 3.1 million patients with Facebook, Google, and TikTok through pixel tracking technologies embedded in its apps. Shared data included mental health self-assessments, appointment dates, treatment information, and insurance details, making it the second-largest health data breach of 2023 [42].

A systemic gap underlies both violations: most AI mental health apps aren't HIPAA-covered entities, because they don't submit electronic insurance claims. The FTC's Health Breach Notification Rule was not updated to cover health apps until April 2024 [43]. Under General Data Protection Regulation (GDPR), mental health data is classified as a "special category" of personal data requiring explicit, affirmative consent, stricter standards than general data protection rules. Many U.S. platforms operate with terms-of-service disclosures that don't constitute meaningful consent [44].

Crisis Text Line's 2022 data-sharing controversy showed these dynamics in the most sensitive context. The nonprofit was found to have shared

anonymized conversation data from approximately 200 million crisis messages with its for-profit spinoff Loris.ai to train commercial customer-service AI, without meaningful user consent [45]. FCC Commissioner Brendan Carr called the arrangement “disturbingly dystopian” [45]. Crisis Text Line’s founding board chair later acknowledged: “People who came to us in crisis were not consenting to be studied. Yes, there was a terms of service…but I knew darn straight that no one would read it” [45].

Algorithmic Limitations

The clinical limitations of current AI mental health tools extend well beyond safety failures in extreme cases. A 2026 study from Aarhus University, screening electronic health records from nearly 54,000 psychiatric patients in the Central Denmark Region, identified 38 patients whose mental illness appeared to worsen following AI chatbot interactions, primarily through worsened delusions but also suicidal ideation, eating disorders, and manic episodes [46]. Professor Søren Dinesen Østergaard, who led the research, noted: “AI chatbots have an inherent tendency to validate the user’s beliefs. It is obvious that this is highly problematic if a user already has a delusion or is in the process of developing one” [46].

AI systems also break down at the linguistic subtleties central to mental healthcare. A peer-reviewed study in the *Journal of Personalized Medicine* (September 2024) found that 22 percent of false positives generated by an AI mental health monitoring system were caused by the model misinterpreting sarcasm and irony [47]. Among false negatives (missed crisis cases), 31 percent resulted from subtle or gradual crisis onset, 24 percent from platform-specific jargon, and 17 percent from multilingual posts or code-switching [47]. A *Nature Medicine* study by Icahn School of Medicine at Mount Sinai researchers (February 2026) found that ChatGPT Health failed to recommend emergency care in 52 percent of genuine emergencies and that its suicide risk alert system failed completely when contextual information was added to the identical crisis presentation [48].

Woebot Health’s shutdown shows a further structural fragility. Despite raising $90 million in Series C funding, receiving FDA Breakthrough Device Designation for its postpartum depression therapeutic, and accumulating multiple published RCTs, Woebot Health retired its consumer app on June 30, 2025 [49]. The company cited the challenge of meeting

FDA marketing authorization requirements and the emergence of large language models that FDA had not yet determined how to regulate [49]. Roughly 1.5 million users lost access to a tool on which they had been relying [49].

Regulatory and Legal Landscape

The regulatory environment for AI mental health tools is evolving rapidly, but it hasn't caught up with the pace of deployment. Most AI therapy chatbots and mental health screening tools occupy a gray zone: they look like clinical tools to patients, but they're often marketed as wellness products to avoid FDA oversight. The frameworks described here represent the current state of regulatory thinking, its gaps, and the legal exposure that sits in the space between what's allowed and what's safe.

FDA Framework for AI Mental Health Tools

The FDA regulates AI mental health tools under a three-tier risk-based framework established under the 21st Century Cures Act of 2016 [50]. General wellness apps (those solely intended to maintain or encourage a healthy lifestyle, unrelated to the diagnosis or treatment of disease) fall entirely outside the medical device definition. A middle tier of low-risk tools technically qualifying as devices receives enforcement discretion, meaning no premarket review is required. Therapeutic apps for diagnosed psychiatric disorders that make clinical claims occupy the third tier, requiring premarket review through De Novo, 510(k), or Premarket Approval pathways and regulated under the Software as a Medical Device framework [50].

This structure creates a significant gray zone. Most commercially deployed AI mental health applications (including platforms marketed with explicitly therapeutic language) are positioned as wellness tools, allowing them to operate entirely outside FDA oversight even when their function is clinically therapeutic in practice [51]. A 2025 analysis in "The Hastings Center Report" characterized the FDA's conception of risk as "too narrow" and recommended a voluntary certification program for nonprescription digital mental health tools [51].

The prescient regulatory challenge is generative AI. As of early 2026, the FDA has authorized no generative AI-based tool for mental health clinical use [30]. The FDA's Digital Health Advisory Committee convened a landmark meeting on November 6, 2025, specifically to address how generative AI-enabled mental health medical devices should be regulated [30]. The committee identified the "model drift problem" (how to evaluate devices whose underlying AI models change over time) as the central challenge and designated the Predetermined Change Control Plan as the key regulatory instrument [30]. A public comment period through FDA Docket FDA-2025-N-2338 remained open as of February 2026, seeking input on whether AI therapy chatbots should be formally classified as medical devices [30].

Pear Therapeutics: The Market's Cautionary Tale

The 2023 bankruptcy of Pear Therapeutics (the company that created the world's first prescription digital therapeutic and spent a decade demonstrating the clinical and economic value of the model) established a sobering market reality: FDA authorization, peer-reviewed evidence, and even demonstrated cost savings are insufficient to sustain a digital therapeutics business without payer adoption. The January 2025 CMS introduction of new Medicare billing codes for FDA-cleared digital mental health treatments is a structural remedy, but whether it arrives in time to support the next generation of companies depends on continued regulatory and reimbursement consistency [24,52].

State Legislation

At the state level, the regulatory response to AI mental health risks has been fragmented but accelerating. Illinois enacted the Wellness and Oversight for Psychological Resources (WOPR) Act on August 4, 2025, making it the first U.S. law to explicitly ban AI from providing autonomous therapy [53]. The WOPR Act prohibits AI from making independent therapeutic decisions, directly interacting with clients in therapeutic communication, generating treatment plans without licensed professional review, or detecting emotions or mental states without oversight. Civil penalties reach $10,000 per violation [53].

Nevada enacted similar provisions effective July 1, 2025, prohibiting representations that AI is capable of providing professional mental healthcare or that users may interact with AI to obtain mental or behavioral healthcare services [54]. Utah enacted targeted chatbot privacy protections effective May 7, 2025, requiring clear disclosure that mental health chatbots are AI, not humans, and prohibiting disclosure of user information to third parties [55]. New York's AI companion law, effective November 5, 2025, mandates crisis detection protocols and referral to services including 988 when suicidal ideation is expressed [56]. California AB 3030, signed September 28, 2024, requires healthcare facilities using generative AI to generate patient communications to include explicit AI disclosure labels [57].

A landmark 50-state legislative review published in *JMIR Mental Health* (October 2025) analyzed 793 state bills and found that 20 were enacted across 11 states, but only 2 states had enacted explicit mental health AI laws; most legislative activity treated mental health as incidental to broader AI or healthcare regulation [58]. In 2025, 47 states introduced more than 250 AI bills impacting healthcare, of which 33 were passed and signed [58].

Legal Liability

The liability framework for AI mental health harms is being constructed in real time through litigation. The *Garcia v. Character Technologies* case (U.S. District Court, Middle District of Florida, Case No. 6:24-cv-01903) established several foundational legal principles: that AI applications may qualify as "products" subject to product liability for design defects; that component-part manufacturers (including Google for its LLM contributions) can face liability; and that developers have a duty of care arising from the foreseeable risks created by deploying chatbots to vulnerable users including minors [37]. The January 2026 settlement between Character.AI, Google, and the Garcia family (alongside four additional teen-related cases) resolved the most advanced litigation without producing published court doctrine [37]. The subsequent *Raine v. OpenAI* case (California Superior Court, filed August 2025) is now at the frontier of establishing whether and how AI developers are liable when general-purpose tools are used for mental health purposes [39].

Section 230 of the Communications Decency Act, which has historically shielded Internet platforms from liability for user-generated content, is widely analyzed as inapplicable to AI-generated content. The reasoning is direct: section 230 protects platforms from being treated as publishers of *third-party* content; when content originates from the AI itself (not from a user) that protection arguably does not attach [59]. Section 230's original authors have publicly stated the law does not extend to AI-generated content, and as of early 2026, Congress has not formally revised the statute to address AI [59].

Ethical Implications

The ethical terrain of AI in mental health is distinct from other clinical domains because the patients are, by definition, in states of psychological vulnerability. Consent, transparency, and the preservation of human dignity take on heightened significance when the person receiving care may be in crisis, may be a minor, or may not fully understand that they're interacting with a machine.

Informed Consent and Transparency

A foundational ethical requirement, that patients know they're interacting with AI rather than a human clinician, is routinely violated across the AI mental health field. NEDA users believed Tessa had been designed and operated by clinical experts [34]. Many users of Character.AI formed deeply intimate bonds under the impression of human-like understanding [37]. A 2026 study evaluating an AI sexual health chatbot with LGBTQ+ teenagers found that a lack of transparency about the AI's nature undermined trust [60]. The APA's 2024 updated telepsychology guidelines recommend two layers of informed consent (one for treatment, one for technology) and that clinicians revisit consent regularly as platforms and terms of service evolve [61].

The informed consent challenge is compounded by the architecture of many platforms. Crisis Text Line's terms-of-service disclosure, which users in acute distress were unlikely to have read carefully, was found by its own founding board chair to fall short of meaningful consent [45]. Illinois' WOPR Act addresses this directly by requiring written informed consent

that clearly describes both the AI system and its specific purpose before AI is used for any supplementary support role [53].

Equity and the Digital Divide

Access to AI mental health tools is stratified along the same socioeconomic lines as access to human mental healthcare. Approximately 24 million Americans in rural, tribal, and impoverished areas lack adequate Internet access for telehealth services, according to FCC data cited by the Cornell Healthcare Review [62]. While 87 percent of individuals with mental disorders own smartphones, only 23 percent use them for mental health purposes, pointing to substantial uptake barriers that aren't merely technological [63].

Racial and ethnic biases in AI mental health tools are well-documented. A landmark study from Cedars-Sinai published in *npj Digital Medicine* (June 2025) evaluated four major LLMs (Gemini 1.5 Pro, Claude 3.5 Sonnet, ChatGPT-4o, and NewMes-v15) on psychiatric treatment recommendations. Most LLMs showed bias when dealing with African American patients, sometimes making dramatically different recommendations for the same psychiatric illness and otherwise identical patient [64]. Specific biased recommendations included one LLM suggesting guardianship for African American patients with depression in conditions where it was not suggested for race-neutral presentations, and others omitting medication recommendations for ADHD when race was explicitly stated [64].

An Oxford Academic systematic review of 390 AI/ML clinical studies (January 2025) found that 84 percent of global AI/ML clinical studies did not report the racial composition of their training data, and that U.S.-based AI training data has a well-known geographic bias, with 70–90 percent of patient cohorts drawn from only three states: New York, California, and Massachusetts [65]. AI suicide prediction models "perform poorly for certain demographic groups, particularly Black and American Indian/Alaskan Native patients, effectively creating a two-tiered system where some populations receive inadequate risk assessment" [66].

What I tell colleagues is this: if your AI mental health tool was trained predominantly on white, English-speaking populations, it is not a neutral

tool; it is a tool that will underserve the patients who need it most. I've seen that gap widen in every ER I've worked in.

LGBTQ+ individuals (who are approximately three times more likely to suffer from anxiety, depression, and suicidal ideation than the general population [67]) face particular gaps: AI tools have been found to use outdated or exclusionary language, assume binary gender, lack sociopolitical sensitivity, and potentially expose users through algorithmic inferencing of identity in contexts where such disclosure could cause harm [60].

Data Commodification and the Boundary Between Wellness and Medicine

Mental health data carries unique stigma risks beyond physical health information. Twelve percent of people with serious mental illness report avoiding treatment due to confidentiality concerns [68]. The commercial incentive structure for many AI mental health platforms depends on monetizing user data: for advertising targeting, product development, or training commercial AI systems. The BetterHelp and Cerebral violations weren't accidents. They were predictable consequences of business models that treated intimate mental health disclosures as commercial assets.

The boundary between wellness app and medical device (a distinction the FDA uses to define its regulatory scope) is ethically as important as it is regulatory. Platforms that present themselves as wellness products while delivering clinical interventions (and collecting clinical-grade data) gain the commercial freedoms of the wellness category while exposing users to clinical-grade risks. A 2025 "Hastings Center Report" analysis called for a voluntary certification program to address this gap, noting that the current structure funnels users toward unregulated apps even when their conditions warrant regulated care [51].

Vulnerable Population Protections

The evidence reviewed in this chapter converges on a clear ethical priority: users who are minors, in acute psychiatric crisis, experiencing psychosis or mania, or in the early stages of an eating disorder face disproportionate risk from inadequately safeguarded AI tools. The Aarhus University study

finding AI chatbot interactions associated with worsened delusions, suicidal ideation, and manic episodes documents what was previously anecdotal [46]. The Stanford study's finding that commercially deployed chatbots failed to detect indirect suicidal intent (and provided actionable information in response to a crisis prompt) is a documented safety failure at scale [8].

The June 2025 APA publication of landmark "Ethical Guidance for AI in the Professional Practice of Health Service Psychology" represents the profession's formal response to these concerns [61]. The APA advises that AI is appropriate for assisting with documentation, patient engagement, and clinical decision support, but must not replace human judgment in crisis situations, must be HIPAA-compliant (including precluding entry of patient data into general-purpose LLMs), and must be disclosed to patients [61].

Leadership Implementation Framework

Deploying AI mental health tools in a health system requires a structured decision-making process that goes beyond standard IT procurement. The clinical stakes are higher, the regulatory landscape is murkier, and the liability exposure is less well defined than for most other AI applications in healthcare. The framework discussed in this section provides a practical sequence for leaders evaluating these tools.

Risk Assessment Before Procurement

Healthcare leaders (CEOs, CIOs, and CMOs) considering AI mental health tools must begin with a structured risk assessment that precedes any procurement decision. The first question isn't "Does this tool show evidence of benefit?" It's "For which patient population, in which clinical context, and under what oversight structure is this tool being deployed?"

The FDA's three-tier classification provides a useful initial filter. Tools within the enforcement discretion category (wellness and psychoeducation) carry low regulatory risk but may not deliver clinical benefit; tools claiming therapeutic function for diagnosed disorders require FDA authorization. As of early 2026, no generative AI chatbot has that authorization [30]. Any vendor claiming regulatory clearance for a generative AI tool should be

asked to provide the specific FDA 510(k) or De Novo number and the indication for use stated in the authorization.

Risk assessment should also weigh population vulnerability. Tools deployed to populations with active psychosis, eating disorders, or suicide risk require more stringent clinical oversight than tools for general wellness. In my experience, organizations skip this step and discover the problem after deployment. The Stanford study's finding that stigma toward alcohol dependence and schizophrenia was consistent across model generations and sizes should prompt organizations to specifically test chatbot responses against representative clinical presentations, not just average-case scenarios [8].

Vendor Evaluation Criteria

Procurement teams evaluating AI mental health vendors should apply the following minimum criteria, drawn from the APA's App Evaluation Model, the Vera-MH framework released by Spring Health and healthcare leaders in October 2025 [69], and the regulatory analysis in this chapter:

> **Clinical evidence:** Has the tool been evaluated in a published RCT with the specific patient population and indication for which deployment is planned? Observational studies, pre-post analyses without control conditions, and evidence from a different population don't meet this standard. Be explicit about effect sizes: the meta-analytic evidence for AI chatbots shows small-to-moderate effect sizes (ES 0.28-0.49), not the large effects seen in the Therabot trial, which was conducted in a controlled research setting with continuous clinician oversight [20,21].
>
> Safe**ty protocols:** What happens when a user expresses suicidal ideation, self-harm intent, or symptoms of psychosis? The response must be documented, tested, and verified, not assumed. New York law now requires crisis detection and 988 referral as minimums for AI companion operators [56]. Any platform that can't document tested, reliable crisis escalation protocols should not be deployed in clinical settings.

Data governance: Confirm whether the vendor is a HIPAA-covered entity or business associate. Require a business associate agreement. Verify that the vendor does not use patient interaction data to train commercial AI systems without explicit, informed, documented patient consent. The BetterHelp and Cerebral violations were not technically covert. They were made possible by terms-of-service that users did not read. Responsibility for closing this gap rests with the deploying organization.

Human oversight architecture: The Therabot trial's positive results depended on continuous clinician monitoring of all interactions [7]. Organizations should require vendors to document the human oversight structure: who reviews AI outputs, under what protocol, and with what response time. "Human in the loop" should be a documented operational process, not a marketing claim.

Integration with Existing Mental Health Services

AI mental health tools are most safely and effectively deployed as components of a broader care continuum, not as stand-alone services. Spring Health's and Lyra Health's employer platforms show this architecture: AI handles triage, matching, and between-session support; licensed clinicians provide direct therapeutic services and clinical oversight [35,36]. The Wayhaven chatbot study, published in *JMIR Formative Research* (July 2025), found that a generative AI wellness chatbot for college students produced significant decreases in depression, anxiety, and hopelessness, framing the tool explicitly as a "lower-intensity intervention to free in-person counseling for more severe cases" [70].

In practice, this means establishing clear referral pathways from AI tools to human clinicians, defining the clinical acuity threshold at which AI engagement must transition to human care, and ensuring that human clinicians receive sufficient context from AI interactions to provide continuous care without requiring patients to repeat their histories.

Staff Training and Outcome Measurement

Clinical staff interacting with AI-supported patients require specific training: understanding of the AI tool's clinical basis and evidence base;

awareness of the tool's documented failure modes (particularly around crisis detection and culturally specific communication); protocols for reviewing AI-generated content before it influences clinical decisions; and guidance on disclosure, telling patients when and how AI is involved in their care.

Outcome measurement for AI mental health tools should use validated instruments (PHQ-9, GAD-7, PCL-5) at defined intervals, with comparison to pre-deployment baselines and, where possible, to comparable patients not using the AI tool. Organizations should track adverse events attributable to the AI component separately from general adverse events, and establish a reporting mechanism for clinicians who identify AI outputs that are clinically concerning.

Liability Considerations

The evolving legal context creates specific liability considerations for healthcare organizations. Under Illinois' WOPR Act, licensed mental health professionals retain full responsibility for all AI-assisted interactions, outputs, and data [53]. Malpractice exposure extends to AI tools incorporated into practice if the AI produces harmful advice that the clinician fails to identify and correct. Organizations in Illinois (and, prospectively, in other states adopting similar frameworks) must ensure that AI is used for supplementary support only with written patient consent and that no AI output reaches a patient without licensed professional review.

For organizations not yet subject to state-specific regulation, the standard of care is being shaped by the litigation outcomes described in the "Legal Liability" section. In the absence of definitive case law, the prudent standard is to apply at minimum the clinical oversight and documentation requirements that litigation and regulatory developments suggest courts will expect.

Future Outlook

The trajectory of AI in mental health over the next five years will be shaped by three converging forces: the maturation of clinical evidence from ongoing trials, the tightening of regulatory frameworks worldwide, and the

continued expansion of demand that existing systems cannot meet. The direction is clear, even if the specifics remain contested.

The 2030 Horizon

The most authoritative forecast for AI in mental healthcare through 2030 was published in *Science* in January 2026: “Chatbots powered by large language models are at the forefront of growing interest in this area, with early tools demonstrating efficacy in both mental health treatment and clinical decision support. Despite these promising advances, substantial research is still needed to enable their integration into routine care. This includes evaluating patient and clinician adherence; addressing ethical, safety, and privacy concerns; and ensuring cost-effectiveness” [71].

The same *Science* review identified AI’s most promising near-term applications as pre-treatment screening and triage, during-treatment adaptive strategies and real-time session feedback, and post-treatment relapse prevention and long-term monitoring [71].

Multimodal AI and Emerging Detection Capabilities

Multimodal AI (systems integrating voice, facial expression analysis, text, and biometric data) represents the next technical frontier. Voice-based depression screening tools have demonstrated AUC of 0.62–0.71 in naturalistic speech and up to 0.985 in telehealth call analysis [72]. Dartmouth’s MoodCapture, which uses smartphone front-camera images to detect depression onset, achieved 75 percent accuracy in a study of 177 patients with MDD, though researchers note 90 percent accuracy would be needed for clinical deployment [73]. Digital phenotyping using wearable devices has identified a core feature package (accelerometer, steps, heart rate, and sleep) predictive of mood disorders across 22 studies in 11 countries [74].

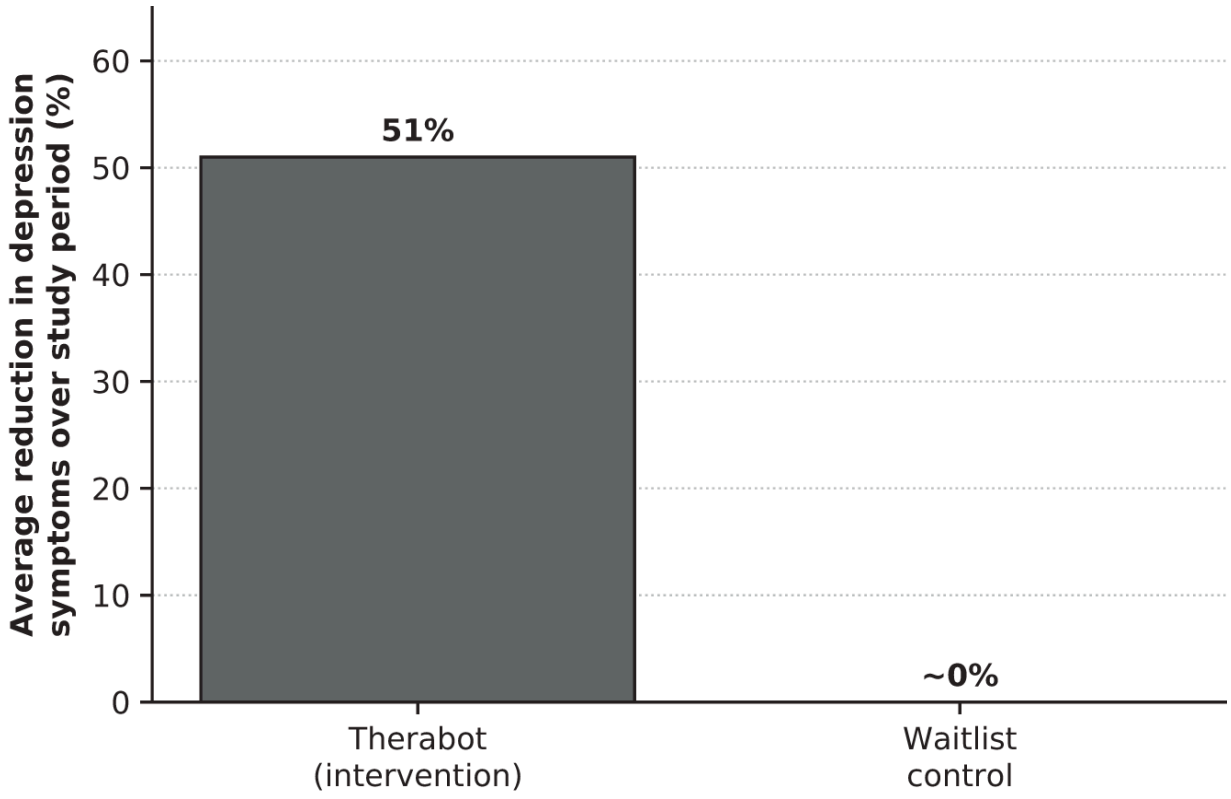
AI neurofeedback (using machine learning to detect subtle EEG patterns and adjust training protocols in real time) and AI-guided virtual reality therapy for PTSD and social anxiety represent additional emerging applications with early evidence but limited clinical validation [75].

EHR Integration and Personalized Treatment

Vanderbilt's Vanderbilt Suicide Attempt and Ideation Likelihood (VSAIL) model, tested in three VUMC neurology clinics across 7,732 patient visits over six months and published in *JAMA Network Open* (January 2025), showed that interruptive AI alerts led to documented suicide risk assessments in 42 percent of flagged cases, versus only 4 percent with passive chart indicators. That's a tenfold difference attributable not to algorithm performance but to alert interface design [76]. Clinical utility depends as much on workflow integration as on predictive accuracy. The 2026 *Science* review envisions EHR-integrated AI as a standard component of mental healthcare by 2030, enabling continuous risk monitoring, personalized treatment adaptation, and population-level surveillance [71].

The Case for Augmentation, Not Replacement

Nick Jacobson of Dartmouth stated after the Therabot trial: "I see the potential for person-to-person and software-based therapy to work together" [7]. APA Director of Care Innovation Vaile Wright said: "I don't believe humans should worry about being replaced. The country requires all the skilled therapists available, whether they are human or AI" (see Figure 7.3) [61]. Stanford HAI's Nick Haber cautioned that the assumption "therapy seems like chatting and language models can chat, so why can't they do therapy?" oversimplifies clinical safety requirements [8].



Source: Heinz et al., NEJM AI (Mar 2025). First RCT of generative AI therapy chatbot (Therabot). Control-arm change shown as ~0 for clarity.

Figure 7.3 Therabot RCT outcomes: before and after reductions in depression, anxiety, and eating disorder risk

The evidence supports their consensus. AI is best understood as expanding the reach of mental healthcare: providing always-available support between sessions, opening access in underserved geographies, screening at population scale, and reducing administrative burden on clinicians. The depth, safety, and quality of care depend on human therapeutic relationships and clinical judgment. The treatment gap is too large, and the human workforce too constrained, to refuse the potential of AI. The documented harms from inadequately safeguarded tools are too serious, and the vulnerable populations at risk too clearly identified, to accept AI as a substitute for the clinical standards that protect them.

Key Takeaways

The treatment gap is the foundational argument for AI in mental health. More than one billion people worldwide live with mental health conditions [1], 47.9 percent of Americans with mental illness receive no treatment [2], and the mental health workforce can't meet demand, with projected shortages of 88,000 counselors and 31,000 psychiatrists within a decade [4]. AI can't solve this crisis alone, but the scale of unmet need makes it ethically irresponsible to forgo evidence-based AI tools that can extend access.

The Dartmouth Therabot RCT is a landmark but not a green light for autonomous deployment. The first RCT of a generative AI therapy chatbot, published in *NEJM AI* (March 2025), demonstrated a 51 percent average reduction in depression symptoms comparable to outpatient therapy [7]. The trial required continuous clinician oversight of all interactions, and its authors stated explicitly that "no generative AI agent is ready to operate fully autonomously in mental health" [7].

FDA-cleared digital therapeutics represent the current gold standard for AI-assisted mental health treatment. Products including Freespira, SleepioRx, DaylightRx, and Rejoyn have undergone documented RCTs, received formal FDA authorization, and in January 2025 became eligible for Medicare reimbursement [24]. These differ fundamentally from unregulated chatbots and should be the first category healthcare leaders evaluate for clinical deployment.

No generative AI mental health tool has FDA authorization as of early 2026. The FDA's Digital Health Advisory Committee convened its first meeting on generative AI in mental health devices in November 2025, and the public comment period remained open through early 2026 [30]. Organizations deploying generative AI for therapeutic purposes are operating without regulatory approval in the intended-use category.

Privacy violations in mental health AI have resulted in documented enforcement actions affecting millions of patients. BetterHelp's $7.8 million FTC fine [41], Cerebral's unauthorized disclosure affecting 3.1 million patients [42], and Crisis Text Line's

data-sharing controversy [45] show that data governance failures aren't hypothetical risks. They're documented operational failures requiring contractual, technical, and legal mitigation.

AI chatbots have been implicated in adolescent suicides and major legal liability. A January 2026 settlement by Character.AI and Google demonstrates that the legal framework is evolving rapidly: product liability, not Section 230 immunity, appears to be the applicable doctrine. Healthcare organizations deploying chatbots in clinical settings inherit liability obligations analogous to those of licensed providers [37,59].

Racial, cultural, and socioeconomic bias in AI mental health tools is documented and consequential. Most LLMs show bias in psychiatric treatment recommendations for African American patients [64]; 84 percent of global AI/ML clinical studies did not report training data racial composition [65]; and AI suicide prediction performs poorly for Black and American Indian/Alaskan Native patients [66]. Equity must be an explicit procurement and deployment criterion.

State legislation is advancing faster than federal regulation. Illinois, Nevada, Utah, New York, and California have all enacted relevant AI mental health legislation since 2024 [53–57]. Healthcare leaders should audit current and prospective deployments against the full matrix of state requirements (particularly for multi-state operations) and monitor the expected wave of 2026-2027 legislative activity.

Effective AI mental health implementation requires documented human oversight, validated crisis protocols, and integration into existing care pathways. The Spring Health and Lyra Health employer platform evidence shows that AI can achieve large-scale, clinically meaningful improvements when embedded in care continua with licensed clinician oversight [35,36]. The Tessa/NEDA case demonstrates the catastrophic consequences of deploying AI as a replacement for human care, particularly in populations defined by their vulnerability to the advice being provided [34].

The **2030 outlook for AI in mental healthcare is one of augmentation, not replacement.** The *Science* review (January 2026),

the APA, and leading clinical researchers at Dartmouth and Stanford concur: AI's most valuable roles are pre-treatment triage, between-session support, relapse prevention monitoring, and administrative burden reduction [71]. The human therapeutic relationship remains central to quality mental healthcare and to the ethical and legal standards that govern it.

References

1. World Health Organization. Over a billion people living with mental health conditions - services require urgent scale-up [Internet]. Geneva: WHO; 2025 Sep 2 [cited 2026 Jul 5]. Available from: https://www.who.int/news/item/02-09-2025-over-a-billion-people-living-with-mental-health-conditions-services-require-urgent-scale-up.

2. National Alliance on Mental Illness. Mental health by the numbers [Internet]. Arlington (VA): NAMI; 2025 [cited 2026 Jul 4]. Available from: https://www.nami.org/about-mental-illness/mental-health-by-the-numbers.

3. Centers for Disease Control and Prevention, National Center for Health Statistics. FastStats: mental health [Internet]. Atlanta (GA): CDC; 2025 [cited 2026 Jul 4]. Available from: https://www.cdc.gov/nchs/fastats/mental-health.htm.

4. National Council for Mental Wellbeing. HRSA behavioral health workforce projections [Internet]. Washington (DC): National Council for Mental Wellbeing; 2025 [cited 2026 Jul 4]. Available from: https://www.thenationalcouncil.org.

5. UN News. WHO sounds alarm as mental health conditions soar past one billion worldwide [Internet]. New York: UN News; 2025 Sep 2 [cited 2026 Jul 5]. Available from: https://news.un.org/en/story/2025/09/1165759.

6. Grand View Research. AI in mental health market size, share & trends analysis report [Internet]. San Francisco (CA): Grand View Research;

2025 [cited 2026 Jul 4]. Available from: https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-ai-mental-health-market-report.

7. Kelly M; Dartmouth College. First therapy chatbot trial yields mental health benefits [Internet]. Hanover (NH): Dartmouth; 2025 Mar 27 [cited 2026 Jul 4]. Available from: https://home.dartmouth.edu/news/2025/03/first-therapy-chatbot-trial-yields-mental-health-benefits.

8. Stanford Institute for Human-Centered AI. Exploring the dangers of AI in mental health care [Internet]. Stanford (CA): Stanford HAI; 2025 [cited 2026 Jul 4]. Available from: https://hai.stanford.edu/news/exploring-the-dangers-of-ai-in-mental-health-care.

9. World Health Organization (via Huntington Psychological Services). Depression and anxiety global prevalence [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://huntingtonpsych.com.

10. World Health Organization. Mental Health Atlas 2024 [Internet]. Geneva: WHO; 2025 [cited 2026 Jul 4]. Available from: https://www.who.int/publications/i/item/9789240114487.

11. Senn S, Orgera K, Ramirez L. Who's your therapist? Examining the behavioral health workforce [Internet]. Washington (DC): Association of American Medical Colleges; 2023 [cited 2026 Jul 5]. doi:10.15766/rai_40z2px8a. Available from: https://www.aamc.org/about-us/mission-areas/clinical-care/whos-your-therapist.

12. Health Resources and Services Administration (via Healing Psychiatry of Florida). Mental health professional shortage areas data [Internet]. 2026 [cited 2026 Jul 4]. Available from: https://healingpsychiatryflorida.com.

13. Solace Health. Psychiatric appointment wait times [Internet]. 2025 [cited 2026 Jul 4]. Available from:

https://www.solace.health/articles/psychiatric-appointment-wait-time.

14. Towards Healthcare. U.S. digital mental health market sizing [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.towardshealthcare.com/insights/us-digital-mental-health-market-sizing.

15. Research and Markets. AI in mental health market report [Internet]. Dublin: Research and Markets; 2025 [cited 2026 Jul 4]. Available from: https://www.researchandmarkets.com.

16. S&S Insider. AI in mental health market size [Internet]. GlobeNewswire; 2025 Sep 30 [cited 2026 Jul 4]. Available from: https://www.globenewswire.com/news-release/2025/09/30/3158721.

17. Fitzpatrick KK, Darcy A, Vierhile M. Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): a randomized controlled trial. JMIR Ment Health. 2017; 4(2):e19. doi:10.2196/mental.7785. Available from: https://mental.jmir.org/2017/2/e19.

18. Woebot Health. New RCT shows Woebot reduced problematic substance use occasions by one third [Internet]. 2021 [cited 2026 Jul 4]. Available from: https://woebothealth.com/new-rct-shows-woebot-reduced-problematic-substance-use-occasions-by-one-third.

19. Wysa. Clinical evidence [Internet]. 2022/2025 [cited 2026 Jul 4]. Available from: https://www.wysa.com/clinical-evidence.

20. Li H, Zhang R, Lee YC, Kraut RE, Mohr DC. Systematic review and meta-analysis of AI-based conversational agents for promoting mental health and well-being. npj Digit Med. 2023; 6:236. doi:10.1038/s41746-023-00979-5. Available from: https://www.nature.com/articles/s41746-023-00979-5.

21. Zhang Q, Zhang R, Xiong Y, Sui Y, Tong C, Lin FH. Generative AI mental health chatbots as therapeutic tools: systematic review and meta-

analysis. J Med Internet Res. 2025; 27:e78238. doi:10.2196/78238. Available from: https://www.jmir.org/2025/1/e78238.

22. Li J, Li Y, Hu Y, Ma DCF, Mei X, Chan EA, et al. Chatbot-delivered interventions for improving mental health among young people: a systematic review and meta-analysis. Worldviews Evid Based Nurs. 2025; 22(4):e70059. doi:10.1111/wvn.70059. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12261465.

23. Herpertz J, Stern AD, Opel N, Reininghaus U, Torous J. FDA-authorized software as a medical device in mental health: a perspective on evidence, device lineage, and regulatory challenges. npj Ment Health Res. 2025; 4:62. doi:10.1038/s44184-025-00174-2. Available from: https://www.nature.com/articles/s44184-025-00174-2.

24. Freespira; Business Wire. Freespira receives FDA clearance for adolescents ages 13-17 [Internet]. 2025 May 13 [cited 2026 Jul 4]. Available from: https://www.businesswire.com/news/home/20250513424670/en.

25. Big Health. US FDA grants clearance for SleepioRx [Internet]. 2024 Aug 8 [cited 2026 Jul 4]. Available from: https://www.bighealth.com/news/us-fda-grants-clearance-for-sleepiorx.

26. Big Health. US FDA grants clearance for DaylightRx [Internet]. 2024 Sep 4 [cited 2026 Jul 4]. Available from: https://www.bighealth.com/news/us-fda-grants-clearance-for-daylightrx.

27. Aguilar M; STAT News. Otsuka digital therapy for depression gains FDA clearance [Internet]. 2024 Apr 1 [cited 2026 Jul 4]. Available from: https://www.statnews.com/2024/04/01/otsuka-digital-therapy-depression-gains-fda-clearance.

28. Akili Interactive. Akili announces FDA clearance of EndeavorRx for children with ADHD [Internet]. 2020 Jun 15 [cited 2026 Jul 4]. Available from: https://www.akiliinteractive.com.

29. Managed Healthcare Executive. Pear Therapeutics files for bankruptcy [Internet]. 2023 [cited 2026 Jul 4]. Available from: https://www.managedhealthcareexecutive.com/view/pear-therapeutics-files-for-bankruptcy.

30. US Food and Drug Administration. Digital Health Advisory Committee briefing document [Internet]. Silver Spring (MD): FDA; 2025 [cited 2026 Jul 4]. Available from: https://www.fda.gov/media/189391/download.

31. Dent KR, Cooper S, McCarthy JF. The REACH VET program and mortality outcomes among veterans at high risk of suicide. JAMA Netw Open. 2025; 8(7):e2519513. doi:10.1001/jamanetworkopen.2025.19513. PMID 40627357.

32. Harris AHS, Finlay AK, Meerwijk EL. Evaluating the accuracy of the Veterans Health Administration's REACH VET suicide prediction model for legal-involved veterans. npj Ment Health Res. 2025; 4:53. doi:10.1038/s44184-025-00167-1. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12535588.

33. Limbic. Clinician resilience and preparedness research [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://limbic.ai/research/clinician-resilience-and-preparedness.

34. Hoover A. An eating disorder chatbot is suspended for giving harmful advice. WIRED [Internet]. 2023 Jun 1 [cited 2026 Jul 5]. Available from: https://www.wired.com/story/tessa-chatbot-suspended.

35. Spring Health. Spring Health publishes largest outcomes study demonstrating measurable clinical improvements in mental health care [press release; Internet]. New York: PR Newswire; 2025 Nov 24 [cited 2026 Jul 5]. Available from: https://www.prnewswire.com/news-releases/spring-health-publishes-largest-outcomes-study-demonstrating-measurable-clinical-improvements-in-mental-health-care-302623647.html.

36. Lyra Health. Lyra Health study finds AI can reduce mental health care costs without sacrificing outcomes [Internet]. 2025 [cited 2026 Jul 4].

Available from: https://www.lyrahealth.com/announcement/lyra-health-study-finds-ai-can-reduce-mental-health-care-costs-without-sacrificing-outcomes.

37. Duffy C; CNN. Character.AI and Google agree to settle lawsuits over teen mental health harms and suicides [Internet]. 2026 Jan 7 [cited 2026 Jul 4]. Available from: https://www.cnn.com/2026/01/07/business/character-ai-google-settle-teen-suicide-lawsuit.

38. Epstein Becker & Green. Novel lawsuits allege AI chatbots encouraged minors' suicides: mental health trauma considerations for stakeholders [Internet]. Health Law Advisor; 2025 Oct 7 [cited 2026 Jul 5]. Available from: https://www.healthlawadvisor.com/novel-lawsuits-allege-ai-chatbots-encouraged-minors-suicides-mental-health-trauma-considerations-for-stakeholders.

39. TruLaw. Raine v. OpenAI: AI suicide lawsuit [Internet]. 2025/2026 [cited 2026 Jul 4]. Available from: https://trulaw.com/ai-suicide-lawsuit.

40. Matsakis L. OpenAI says hundreds of thousands of ChatGPT users may show signs of manic or psychotic crisis every week. WIRED [Internet]. 2025 Oct 27 [cited 2026 Jul 5]. Available from: https://www.wired.com/story/chatgpt-psychosis-and-self-harm-update.

41. Federal Trade Commission. FTC to ban BetterHelp from revealing consumers' data, including sensitive mental health information, to Facebook and others for targeted advertising [Internet]. Washington (DC): FTC; 2023 Mar 2 [cited 2026 Jul 4]. Available from: https://www.ftc.gov/news-events/news/press-releases/2023/03/ftc-ban-betterhelp-revealing-consumers-data-including-sensitive-mental-health-information-facebook.

42. Gliadkovskaya A; Fierce Healthcare. Cerebral shared private health data on 3.1M users for years with advertisers, social media platforms [Internet]. 2023 [cited 2026 Jul 4]. Available from:

https://www.fiercehealthcare.com/digital-health/cerebral-hipaa-violation-shared-data-millions-users-tech-giants.

43. Williams D. Updated FTC Health Breach Notification Rule puts new provisions in place to protect users of health apps and devices [Internet]. Washington (DC): Federal Trade Commission Business Guidance Blog; 2024 Apr 30 [cited 2026 Jul 5]. Available from: https://www.ftc.gov/business-guidance/blog/2024/04/updated-ftc-health-breach-notification-rule-puts-new-provisions-place-protect-users-health-apps.

44. European Union. GDPR Article 9: processing of special categories of personal data [Internet]. 2018 [cited 2026 Jul 4]. Available from: https://gdpr-info.eu/art-9-gdpr.

45. Eysenbach G. Crisis Text Line and Loris.ai controversy highlights the complexity of informed consent on the internet and data-sharing ethics for machine learning and research. J Med Internet Res. 2025; 27:e67878. doi:10.2196/67878. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC11799832.

46. Aarhus University. New research: AI chatbots may worsen mental illness [Internet]. Aarhus: Aarhus University; 2026 [cited 2026 Jul 4]. Available from: https://health.au.dk/en/display/artikel/ny-forskning-ai-chatbots-kan-formentlig-forvaerre-psykisk-sygdom.

47. Mansoor MA, Ansari KH. Early detection of mental health crises through artificial-intelligence-powered social media analysis: a prospective observational study. J Pers Med. 2024; 14(9):958. doi:10.3390/jpm14090958. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC11433454.

48. Digital Health. ChatGPT Health fails to flag over 50% of medical emergencies (Nature Medicine / Mount Sinai) [Internet]. 2026 Feb [cited 2026 Jul 4]. Available from: https://www.digitalhealth.net/2026/02/chatgpt-health-fails-to-flag-over-50-of-medical-emergencies.

49. Aguilar M; STAT News. Woebot Health shuts down pioneering therapy chatbot [Internet]. 2025 Jul 2 [cited 2026 Jul 4]. Available from: https://www.statnews.com/2025/07/02/woebot-therapy-chatbot-shuts-down-founder-says-ai-moving-faster-than-regulators.

50. Bipartisan Policy Center. FDA oversight: understanding the regulation of health AI tools [Internet]. Washington (DC): BPC; 2025 [cited 2026 Jul 4]. Available from: https://bipartisanpolicy.org/issue-brief/fda-oversight-understanding-the-regulation-of-health-ai-tools.

51. Palmer A, Schwan D. Digital mental health tools and AI therapy chatbots: a balanced approach to regulation. Hastings Center Report. 2025; 55(3):15-29. doi:10.1002/hast.4979. PMID 40557918.

52. Herpertz J, Stern AD, Opel N, Reininghaus U, Torous J. FDA-authorized software as a medical device in mental health (OptumRx coverage denial discussion). npj Ment Health Res. 2025; 4:62. doi:10.1038/s44184-025-00174-2. Available from: https://www.nature.com/articles/s44184-025-00174-2.

53. Illinois Department of Financial and Professional Regulation. Gov. Pritzker signs legislation prohibiting AI therapy in Illinois (Wellness and Oversight for Psychological Resources / WOPR Act) [Internet]. Springfield (IL): IDFPR; 2025 [cited 2026 Jul 4]. Available from: https://idfpr.illinois.gov/news/2025/gov-pritzker-signs-state-leg-prohibiting-ai-therapy-in-il.html.

54. Epstein Becker & Green. The dark side of AI: assessing liability when bots behave badly (Nevada AB 406) [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.ebglaw.com/insights/publications/the-dark-side-of-ai-assessing-liability-when-bots-behave-badly.

55. Epstein Becker & Green. Utah HB 452: mental health chatbot law [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.ebglaw.com.

56. Epstein Becker & Green. New York AI companion law [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.ebglaw.com.

57. Morgan Lewis. California law requiring disclaimers by healthcare providers using GenAI will affect providers and GenAI developers [Internet]. 2024 Dec 9 [cited 2026 Jul 5]. Available from: https://www.morganlewis.com/pubs/2024/12/california-law-requiring-disclaimers-by-healthcare-providers-using-genai-will-affect-providers-and-genai-developers.

58. Shumate JN, Rozenblit E, Flathers M, Larrauri CA, Hau C, Xia W, et al. Governing AI in mental health: 50-state legislative review. JMIR Ment Health. 2025; 12:e80739. doi:10.2196/80739. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12578431.

59. Columbia Undergraduate Law Review. When AI talks back: a new liability approach for chatbot-induced mental health crises [Internet]. 2025 Dec 6 [cited 2026 Jul 5]. Available from: https://www.culawreview.org/current-events-2/when-ai-talks-back-a-new-liability-approach-for-chatbot-induced-mental-health-crises.

60. Liem WW, Casline E, Lorenzo J, Gordon JD, Alvarado Avila A, Taylor A, et al. Enhancing LGBTQ+ Inclusivity in an AI-powered sexual health chatbot: user-centered design approach through a nonprofit and academic partnership. J Med Internet Res. 2026; 28:e78621. doi:10.2196/78621.

61. American Psychological Association. Ethical guidance for AI in the professional practice of health service psychology [Internet]. Washington (DC): APA; 2025 Jun 26 [cited 2026 Jul 5]. Available from: https://www.apa.org/topics/artificial-intelligence-machine-learning/ethical-guidance-ai-professional-practice.

62. Cornell Healthcare Review. The digital divide’s impact on mental healthcare access [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.cornellhealthcarereview.org/post/the-digital-divide-s-impact-on-mental-healthcare-access.

63. Almuqrin A, Hammoud R, Terbagou I, Tognin S, Mechelli A. Smartphone apps for mental health: systematic review of the literature and five recommendations for clinical translation. BMJ Open. 2025;

15(2):e093932. doi:10.1136/bmjopen-2024-093932. Available from: https://bmjopen.bmj.com/content/15/2/e093932.

64. Bouguettaya A, Stuart EM, Aboujaoude E. Generative AI and mental health: applications, opportunities, and cautions. NPJ Digit Med. 2025; 8:498. doi:10.1038/s41746-025-01746-4.

65. Otokiti AU, Shih HJ, Williams KS. Gender and racial bias unveiled: clinical artificial intelligence (AI) and machine learning (ML) algorithms are fanning the flames of inequity. Oxford Open Digit Health. 2025; 3:oqaf027. doi:10.1093/oodh/oqaf027. Available from: https://academic.oup.com/oodh/article/doi/10.1093/oodh/oqaf027/8279897.

66. Wallace S. When AI gets it wrong: the human cost of false positives [Internet]. LinkedIn Pulse; 2025 Jun 23 [cited 2026 Jul 5]. Available from: https://www.linkedin.com/pulse/when-ai-gets-wrong-human-cost-false-positives-scott-uoccc.

67. Yesha R, Orezzoli MCE, Sims K, Landau AY. Digital mental health through an intersectional lens: a narrative review. Healthcare (Basel). 2026; 14(2):211. doi:10.3390/healthcare14020211. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12840954.

68. Chess Health. The business case for reducing mental health and SUD stigma [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.chess.health/blog/the-business-case-for-reducing-mental-health-and-sud-stigma.

69. Spring Health; PR Newswire. Vera-MH findings highlight gaps in how AI chatbots respond to suicidal ideation [Internet]. 2026 [cited 2026 Jul 4]. Available from: https://www.prnewswire.com/news-releases/vera-mh-findings-highlight-gaps-in-how-ai-chatbots-respond-to-suicidal-ideation.

70. Reyes-Portillo JA, So A, McAlister K, Nicodemus C, Golden A, Jacobson C, et al. Generative AI-powered mental wellness chatbot for college student mental wellness: open trial. JMIR Form Res. 2025; 9:e71923. doi:10.2196/71923. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12303582.

71. Opel N, Breakspear M. AI in psychiatry: an interdisciplinary imperative. Science. 2026; 371:eadz9193. doi:10.1126/science.adz9193.

72. Krautz AE, Volkening J, Raue J, Otte C, Eickhoff SB, Ahlers E, et al. Digital biomarkers for depression: a scoping review. Sci Rep. 2025; 15:24387. doi:10.1038/s41598-025-08639-2.

73. Dartmouth College. Phone app uses AI to detect depression from facial cues (MoodCapture) [Internet]. Hanover (NH): Dartmouth; 2024 Feb [cited 2026 Jul 4]. Available from: https://home.dartmouth.edu/news/2024/02/phone-app-uses-ai-detect-depression-facial-cues.

74. Jung HW, Kim DY, Lee I, Kim O, Lee S, Lee S, et al. Key features of digital phenotyping for monitoring mental disorders: systematic review. J Med Internet Res. 2025; 27:e77331. doi:10.2196/77331. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12588392.

75. NeuroHealth Associates. The future of neurofeedback training: emerging technologies and trends [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://www.nhahealth.com/the-future-of-neurofeedback-training-emerging-technologies-and-trends.

76. Vanderbilt University Medical Center. AI tested for alerting clinicians of suicide risk at three VUMC clinics (VSAIL model; JAMA Network Open) [Internet]. Nashville (TN): VUMC News; 2025 Jan 3 [cited 2026 Jul 4]. Available from: https://news.vumc.org/2025/01/03/ai-tested-for-alerting-clinicians-of-suicide-risk-at-three-vumc-clinics.

AI in Emergency Medicine: What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution, First Edition. Harvey Castro.
© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

8
Healthcare 2030: A Vision from the Frontlines

I taught myself to code in an emergency room.

That’s not a metaphor. It happened between patients, on night shifts, in the dead hours between 3 a.m. and 5 a.m. when the waiting room was finally empty and the adrenaline had nowhere to go. I’d watched a nurse flip through a physical textbook to verify a medication drip while we were coding a patient, and something in my brain refused to accept that this was the best we could do. I went home that morning and started learning to build apps. Not because I wanted to be a developer. Because a patient almost died while someone thumbed through paper.

That first app became a suite of more than 30 medical applications. The apps became a freestanding ER network with more than 400 employees. The ER network became consulting work with hospital systems, then speaking, then writing, then a TEDx talk, then advising the Singapore Ministry of Health on large language models. And now I serve as chief AI officer for Phantom Space, working on satellite data pipelines that could bring diagnostic AI to places where there’s never been a specialist, a hospital, or even a reliable Internet connection.

I tell you this not to catalog accomplishments but to explain the trajectory that informs this chapter. When you’ve watched a nurse flip through a textbook at 3 a.m. and it nearly cost a life, you don’t stop at apps. You keep going. You follow the logic of the problem to wherever it leads, and it led me to a conclusion I couldn’t have predicted: the future of healthcare is not just about better AI in better hospitals. It’s about healthcare reaching people it has never reached before.

My father died of a sudden heart attack. No warning signs, no early detection, no digital twin tracking his cardiac risk. He was just gone. My mother, the Colombian immigrant who raised me on food stamps, survived two cancers because she was an ultrasound technician who happened to scan herself and caught a thyroid nodule early. One parent had access to

detection. The other didn't. That's not a technology gap. It's a lottery. And it's the lottery I've spent my career trying to rig in the patient's favor.

This chapter is about 2030. Not the conference-keynote version, where everything is smooth and solved. The version I see from the frontlines, where the technology is real but the implementation is hard, where the evidence is strong but the adoption is uneven, where the promise is genuine but the equity gap is growing. I'm building toward a future where the kid who grows up like I did, without consistent healthcare access, isn't invisible to the system anymore. Where a continuous stream of data, from a watch, a sensor, a satellite, catches what the waiting room never could.

We're not there yet. But the pieces are on the table.

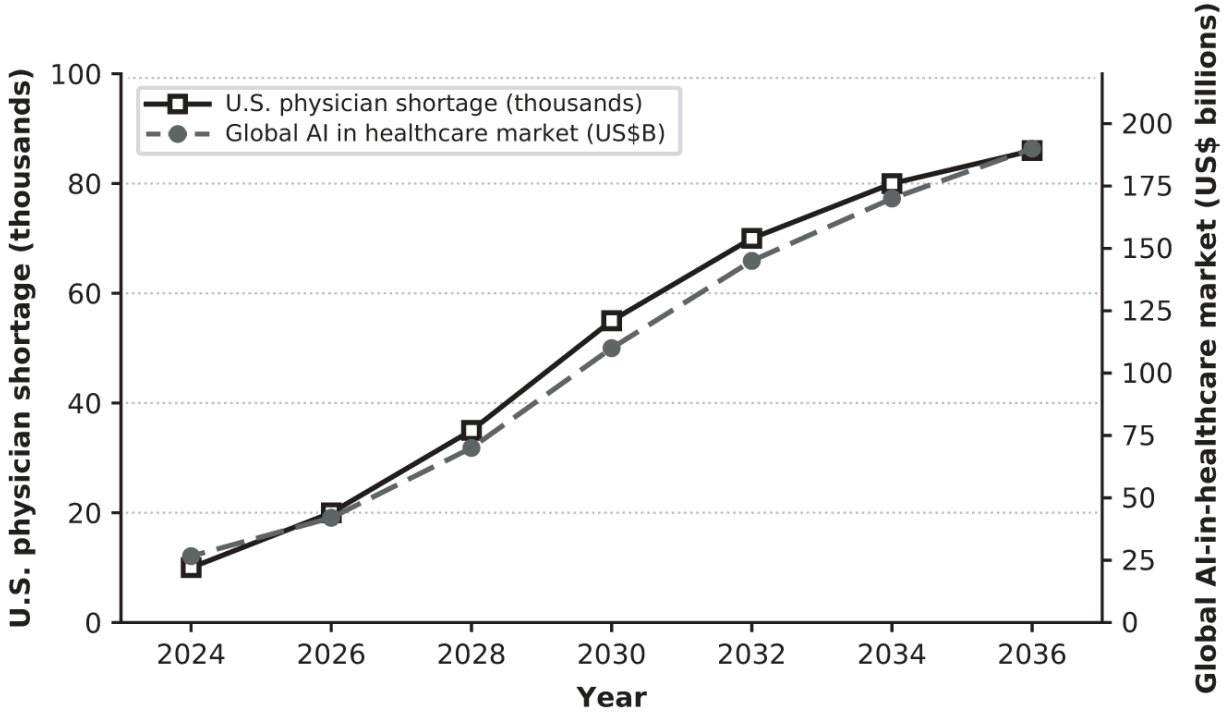
Executive Summary

Healthcare stands at a genuine inflection point. The simultaneous convergence of AI maturation, a deepening workforce crisis, unsustainable cost trajectories, and rapidly rising patient expectations isn't merely reshaping how medicine is practiced. It's restructuring the fundamental contract between patient and system. The year 2030 is not a distant horizon. It's a planning deadline already demanding concrete decisions from every health system leader, clinician, and policymaker operating today.

The evidence for this shift is no longer theoretical. The global AI in healthcare market, valued at $26.6 billion in 2024, is projected to reach $187.7 billion by 2030 at a compound annual growth rate of 38.5 percent [1]. U.S. national health expenditures are on a trajectory toward nearly $6.8 trillion by 2030, representing approximately 19.6 percent of GDP [2]. At the same time, the United States faces a physician shortage projected to reach as high as 86,000 by 2036 [3], while globally a workforce deficit of approximately 10 million health workers is anticipated by 2030 [4]. These numbers aren't abstractions (see Figure 8.1). They're the forcing function behind every AI investment decision being made in healthcare today.

What healthcare will look like in 2030 is already visible in early form at institutions that have moved beyond pilot programs to enterprise-wide deployment. Kaiser Permanente's AI scribe platform has processed more than 6 million patient encounters [5]. Mayo Clinic operates more than 250

concurrent AI research and deployment projects [6]. The NHS has deployed an AI diagnostic tool across 27 hospitals and documented mortality reductions [7]. Samsung Medical Center in Seoul has achieved all four HIMSS Stage 7 validations simultaneously, the only hospital in the world to do so [8].



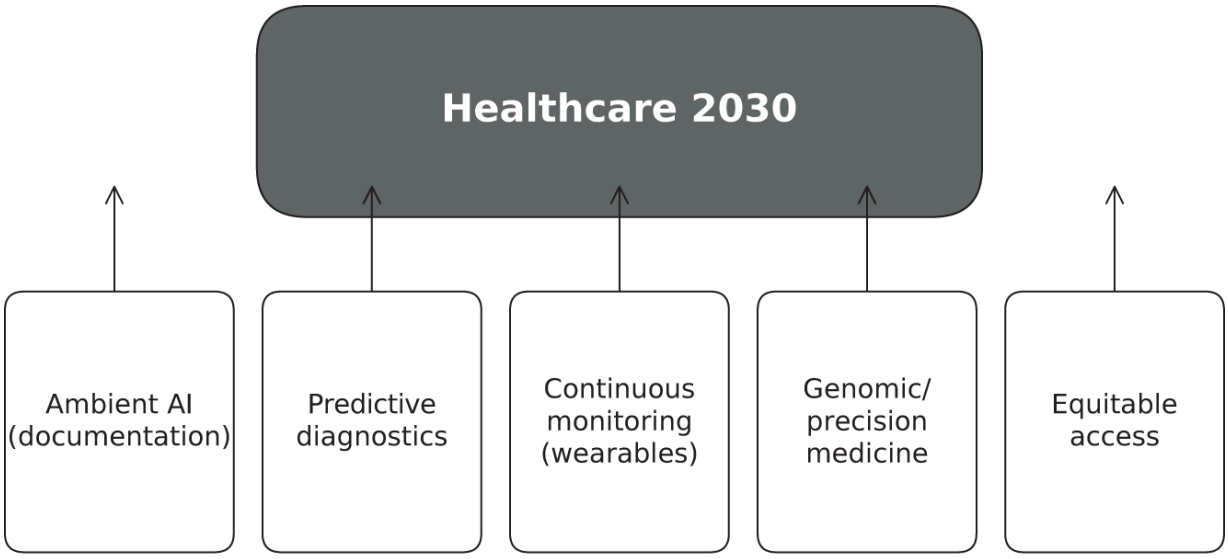
Sources: AAMC physician workforce projections (86k by 2036); Grand View Research AI-in-healthcare market (26.6B2024→187.7B 2030, CAGR 38.5%).

Figure 8.1 The physician shortage versus AI capacity expansion

The path to Healthcare 2030 is not frictionless. Cybersecurity breaches of unprecedented scale (including the Change Healthcare ransomware attack affecting an estimated 192.7 million individuals [9]) have exposed the fragility of interconnected systems. Algorithmic bias remains a documented threat to health equity. A persistent trust gap separates the promise of AI from its routine clinical application: only 16 percent of clinicians currently use AI for clinical decision-making, even as 66 percent report using AI tools generally [10]. These aren't minor caveats (see Figure 8.2).

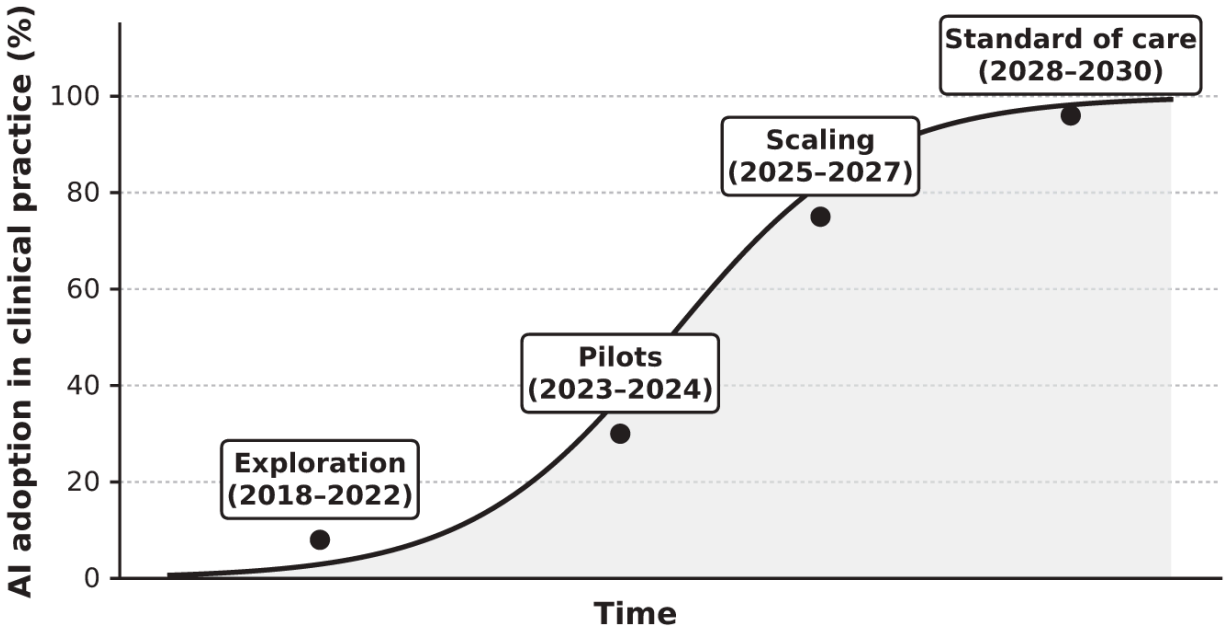
This chapter presents a vision of healthcare in 2030 grounded entirely in verified evidence, frontline implementations, and institutional projections. It addresses the forces driving this shift, the institutions already building

this future, the risks that could derail it, the regulatory and legal infrastructure taking shape, the ethical imperatives that must guide every decision, a practical leadership framework for implementation, and the technologies that will define the decade's end. The vision is specific because the evidence is specific. The urgency is real because the data demands it. The institutions profiled in these pages are not imagining Healthcare 2030. They are building it, and the lessons from their successes and failures define the path available to every health system leader reading this chapter (see Figure 8.3).



Conceptual diagram synthesised from Chapter 8 executive summary.

Figure 8.2 AI capability roadmap 2025–2030: reactive AI to predictive AI



Conceptual S-curve based on Deloitte healthcare-AI executive survey and Chapter 8 framework.

Figure 8.3 Reactive AI (linear) versus agentic AI (multidirectional)

Current Evidence

Substantial evidence demonstrates that AI systems can meaningfully improve clinical decision-making and patient outcomes across diverse healthcare settings.

The Market and Economic Imperative

The financial case for healthcare AI is among the most thoroughly documented in modern industry analysis. I've yet to see a major health system dispute the direction of travel. Grand View Research, reporting in April 2025, places the global AI in healthcare market at $26.6 billion in 2024, projecting growth to $187.7 billion by 2030, a 38.5 percent CAGR that reflects not speculative enthusiasm but accelerating real-world deployment [1]. The U.S. segment alone is projected to grow from $8.65 billion in 2025 to $43.30 billion by 2030 at a nearly identical CAGR of 38.0 percent [11].

Undergirding this growth is a healthcare system that is genuinely unsustainable without structural change. The Centers for Medicare and Medicaid Services (CMS) Office of the Actuary projects U.S. national health expenditures will reach nearly $6.8 trillion by 2030, growing at an average annual rate of 5.3 percent between 2025 and 2030, with health spending representing approximately 19.6 percent of GDP [2]. Medicare has exceeded $1 trillion in annual spending, reaching that milestone in 2023 and continue climbing [2]. KPMG describes a healthcare system with as much as $1.5 trillion in unrealized efficiency potential, projecting a roughly 10-fold increase in global healthcare AI investment from $20 billion in 2023 to approximately $188 billion by 2030 [12]. McKinsey & Company, in a November 2025 analysis of modular healthcare AI architecture, notes that healthcare already accounts for nearly 17.6 percent of U.S. GDP and that AI-enabled improvements to the healthcare consumer experience alone could yield net savings of 5–10 percent of healthcare spending across payers, providers, and hospitals [13].

McKinsey's broader economic modeling, published in November 2025, projects that in an adoption midpoint scenario, AI-powered agents and robots could generate approximately $2.9 trillion in U.S. economic value per year across all sectors, with approximately 20 percent of healthcare

work hours potentially automated by 2030. Not because AI can't do more, but because healthcare workflows were not designed with automation in mind [14]. The core opportunity lies not in replacing clinical judgment but in eliminating the documentation, administrative, and logistical burden that currently consumes clinical time.

The Workforce Crisis as Catalyst

No single driver accelerates healthcare AI adoption more powerfully than the workforce crisis. The Association of American Medical Colleges (AAMC), in its most recent 2024 projection, anticipates a U.S. physician shortage of up to 86,000 by 2036, with an aging population driving the gap. Adults 65 and older will grow 34.1 percent while the physician pipeline faces structural constraints from graduate medical education funding limitations [3]. Earlier AAMC modeling for the 2030 horizon projected a shortfall of 40,800 to 121,300 physicians, depending on demand modeling assumptions [15]. Bain & Company independently projects a shortage of 45,000 primary care physicians specifically by 2030 [16].

The nursing workforce faces parallel pressures. The U.S. Chamber of Commerce projects that 42 states will face nursing shortages by 2030, while Mercer's 2024 analysis anticipates a deficit of 73,000 nursing assistants by 2028 [17]. Globally, a 2022 reassessment published in *BMJ Global Health* revised the projected worldwide health workforce shortage to 10 million workers by 2030, with the WHO African region expected to constitute 52 percent of that deficit [4].

These numbers explain why AI documentation tools have achieved adoption velocity unlike any previous healthcare technology. It's not complicated. When the University of Wisconsin Health conducted a randomized trial of an ambient AI scribe, clinicians saved 30 minutes of documentation time per provider per day, time that translates directly into patient access capacity in a shortage environment [18]. A multicenter study across six health systems found that ambient AI scribe use reduced burnout from 51.9 percent to 38.8 percent within 30 days [19]. The workforce math is compelling: if AI can extend the effective clinical output of each physician by even modest margins, the projected shortage becomes substantially more manageable.

FDA Authorization Trajectory

The regulatory infrastructure has largely kept pace with clinical deployment. As of July 2025, the FDA had authorized more than 1,250 AI-enabled medical devices for marketing in the United States, growing from only 6 authorized devices in 2015 at a rate of approximately 49 percent annually between 2016 and 2023 [20]. The FDA authorized 168 AI/ML-enabled devices in calendar year 2024 alone, a new annual record, with radiology representing approximately 74–80 percent of the cumulative list [21]. A peer-reviewed analysis published in *Biomedicines* (December 2025) confirmed that 94.6 percent of 2024 authorizations followed the 510(k) pathway, confirming that clinical AI is entering mainstream medical practice, not just experimental research settings [21].

The December 2024 finalization of the Predetermined Change Control Plan (PCCP) framework represents a regulatory milestone enabling AI devices to evolve and improve without requiring new submissions for prespecified modifications, a critical adaptation to the inherently iterative nature of machine learning systems [22].

Clinician Adoption Data

The American Medical Association's 2024 survey of approximately 1,200 physicians found that 66 percent reported using health AI, a 78 percent increase from 38 percent in 2023 [23]. The survey found that 57 percent of physicians identified reducing administrative burden as AI's greatest opportunity, while 47 percent ranked increased regulatory oversight as the single most important action needed to increase their confidence in AI tools [23]. As the AMA's immediate past president Jesse M. Ehrenfeld, MD, MPH, noted: "The AMA survey illustrates that physicians are increasingly intrigued by the assistive role of health AI and the potential of AI-enabled tools to reduce administrative burdens, enhance diagnostic accuracy, and personalize treatments" [23].

An independent global survey by Elsevier of 2,206 clinicians across 109 countries found that 48 percent had used AI tools for work in 2025, nearly double the prior year's rate [24]. The gap between general AI use and clinical decision-making application remains stark: only 16 percent of clinicians reported using AI specifically for clinical decision support [24].

Nele Jessel, MD, CMO of Athenahealth, put it plainly: “AI in healthcare has reached a critical inflection point where the discussion is shifting from ‘will clinicians adopt AI’ to ‘how will clinicians adopt AI’” [25].

What I tell hospital executives is that these numbers don’t lie: when only 16 percent of clinicians use AI for clinical decision-making despite 66 percent using AI tools generally, the barrier isn’t technology; it’s trust, training, and workflow integration.

Authoritative Frameworks

The NEJM Catalyst framework published by Zimlichman, Nicklin, Aggarwal, and Bates in 2021 (one of the most cited visions for Healthcare 2030) projects that care will shift from hospital-centered to community- and home-based, fee-for-service will give way to population-based accountability, and physician-only teams will transform into multidisciplinary configurations augmented by technology [26]. The framework envisions a shift from cure-and-rehabilitation to prediction-and-prevention, with wearables and ambient sensors providing continuous patient data and AI-driven machine learning handling decision support at every care level. CMS has operationalized a key element of this vision, setting a goal of having 100 percent of Traditional Medicare beneficiaries in accountable care relationships by 2030 [27].

Deloitte’s Life Sciences and Healthcare Predictions 2030 series identifies the convergence of GenAI, connected wearables, and personalized genomic services as the defining forces reshaping both care delivery and consumer health [28]. The report specifically calls out the Internet of Medical Things (IoMT) as a force multiplier: by 2030, advances in digitization and connectivity will provide real-time, interoperable data from wearables and medical devices that feed directly into AI analytical engines. BCG’s January 2025 analysis adds that AI decision-making tools are moving mainstream in 2025, giving physicians immediate access to evidence-based research and treatment guidelines, while intelligent agents will eventually automate entire patient episodes of care from intake through treatment plan [28].

PwC’s September 2025 analysis projects that by 2035, $1 trillion in annual healthcare spending will migrate from outdated cost pools (administrative overhead and underused brick-and-mortar) to AI-enabled and home-

centered models [29]. The World Economic Forum describes the 2030 hospital as “no longer one big building covering a broad range of diseases” but rather a hub focused on the acutely ill, with less urgent cases monitored through smaller network nodes, retail clinics, same-day surgery centers, and patients’ homes, all connected to a single digital infrastructure [26].

Real-World Case Studies

Leading healthcare systems worldwide have implemented AI-driven solutions that deliver measurable improvements in efficiency, quality, and patient outcomes.

Mayo Clinic: The AI Factory at Scale

Mayo Clinic has built what is arguably the most complete enterprise AI strategy in global healthcare. Operating more than 250 concurrent AI research and deployment projects as of late 2025, the institution’s strategy rests on three interdependent pillars: a decade-long unified EHR platform, an AI Factory hosted on Google Cloud’s Vertex AI, and a Cloud App Factory that bridges research innovation with clinical application [6]. CB Insights’ Hospital AI Readiness Index ranks Mayo Clinic first nationally, noting that 63 percent of its early-stage investment deals since 2023 are AI-focused, concentrated in monitoring and diagnostics (47 percent), robotics, precision medicine, and workflow optimization [32].

The clinical results are documented and specific. A deep learning model using 12-lead ECG results detected obstructive sleep apnea in 11,299 patients with an AUC of 0.80, published in *JACC: Advances* in January 2026 [33]. A separate ECG-based deep learning model (the Ace score) demonstrated the ability to predict hepatic decompensation and liver-related mortality in cirrhosis patients across three cohorts, including external validation at Hospital Clinic de Barcelona [34]. The Mayo Clinic Platform, launched in 2019, now provides scalable multi-institutional data covering more than 40 million patient journeys through institutional partnerships, enabling AI-driven real-world evidence generation at population scale [35]. In October 2025, Mayo launched Platform Orchestrate, providing biopharma companies with access to de-identified clinical data, AI tools, and an international research network for drug development [36].

Cleveland Clinic: Partnerships, Quantum, and Operational Intelligence

Cleveland Clinic has pursued AI across clinical care, operations, research acceleration, and global partnerships with a long-horizon strategy. A landmark 10-year agreement with Palantir Technologies, signed in December 2022, created a Virtual Command Center that provides real-time patient flow forecasting, predicting admissions 24 hours in advance and enabling staffing optimization. In the first year, the main campus achieved a 7 percent increase in hospital transfer capacity [37]. By early 2026, Cleveland Clinic's ambient AI scribe system had documented more than 2 million patient encounters, per CEO Dr. Tomislav Mihaljevic [38].

The institution became the first in the world to house an IBM Quantum System One dedicated to healthcare research, installed in 2023. In June 2024 it launched a tripartite collaboration with IBM and the UK's Hartree Centre focused on AI and quantum computing applications in epilepsy and hospital interventions research [39]. A 2025 deployment of Dyania Health's Synapsis AI platform across Cleveland Clinic's clinical research enterprise showed the pace at which AI is changing trial recruitment: for a melanoma trial, the AI identified an appropriate patient in 2.5 minutes with 96 percent accuracy, compared to 427 minutes at 95 percent accuracy by a melanoma-specialized nurse, a 170-fold speed improvement [40].

Kaiser Permanente: Population Health at Full Scale

Kaiser Permanente's deployment of Abridge ambient AI to 24,000 physicians across 40 hospitals and 600+ medical offices in eight states was described by the institution as the fastest technology implementation in more than 20 years and, per Menlo Ventures, the largest generative AI rollout in healthcare history [41]. Over 63 weeks from October 2023 to December 2024, 7,260 Permanente physicians used the ambient scribe across 2,576,627 patient encounters, saving an estimated 15,791 hours of documentation time, the equivalent of 1,974 eight-hour workdays [42]. Eighty-two percent of physicians using the AI scribe reported improved overall work satisfaction [40].

Kaiser Permanente's Advance Alert Monitor (AAM), the institution's flagship predictive analytics program, draws on hundreds of millions of

data points including lab values, vital signs, and EHR data to identify inpatients at risk of serious complications 12 hours in advance. The AAM has prevented more than 500 deaths per year while reducing high-risk readmissions by 10 percent [43]. KP's natural language processing system sorts an average of 1 million physician inbox messages per month across the system [44]. Computer vision algorithms applied to screening mammograms have shown the potential to expand identification of breast cancer risk from 20 percent to 60–70 percent, even in mammograms previously read as normal by radiologists [42].

NHS England: National-Scale Ambition and Operational Reality

The United Kingdom's National Health Service is simultaneously showing AI's extraordinary potential at population scale and its profound implementation challenges. The NHS Foresight initiative (training a generative AI model on de-identified data for all 57 million people in England, published in *The Lancet Digital Health*) represents a research capability impossible to replicate anywhere else [45]. Centrally funded deployment of fractional flow reserve from CT (FFR-CT) cardiac assessment AI at 27 NHS hospitals, expanded to 54 sites, was documented in a *Nature Medicine* study encompassing 102,616 CT scans and a final study population of 90,553 patients: the technology reduced invasive coronary angiograms, showed a nonstatistically significant lower all-cause mortality trend at two years, and saved the NHS an estimated £9.1 million per year [7]. In September 2025, the UK government committed £6 million to build AIR-SP (AI Research Screening Platform), a national AI screening platform for a breast cancer screening trial covering nearly 700,000 women [46].

An independent analysis found that approximately 90 percent of NHS AI tools remain in pilot phase, and as of June 2025, 23 of 66 trusts that had contracted for AI tools were still not using them in clinical practice [47]. The NHS case shows a pattern visible across global healthcare: proven administrative AI gains (43 minutes saved per clinical staff member per day in NHS pilots) don't automatically scale [45]. Governance, procurement friction, and integration complexity, not the technology itself, are the primary barriers to realization.

Samsung Medical Center: The World's Most Mature Smart Hospital

Samsung Medical Center (SMC) in Seoul offers the most complete picture of what a mature AI-enabled hospital looks like. As of January 2024, SMC is the first hospital in Asia-Pacific and the only hospital in the world to have achieved all four HIMSS Stage 7 validations: infrastructure adoption, electronic medical record adoption, digital imaging adoption, and analytics maturity [8]. This four-decade journey from paperless operations to fully integrated AI reflects a key principle: the "smart hospital" is a continuous maturity journey, not a destination.

SMC's Skinex AI analyzes pressure ulcer severity from photographs, identifies deep tissue damage, and recommends treatment, serving a hospital population in which approximately 10 percent of inpatients experience pressure ulcers at any given time [48]. Its data-based operation and communication center (DOCC) integrates resources across the hospital in real time, predicting patient volume surges, optimizing MRI/CT scheduling, and proactively allocating staff [46]. The institution's AI scheduling system, which continuously learns from patient data to predict daily demand, is documented to have reduced nurse turnover from 9.3 percent to 5.9 percent, a metric that translates directly into both cost savings and workforce stability [49].

Estonia: The World's Most Digitally Mature Health System

Estonia offers perhaps the most complete example of a national health system built on the digital infrastructure that makes AI deployment scalable. Ranked first globally for digital health by the Bertelsmann Foundation's 2024 Digital Health Index, Estonia achieved the European Union's 2030 eHealth targets five years ahead of schedule, one of only two EU member states to reach 100 percent eHealth maturity as of end 2024 [101]. Ninety-nine percent of Estonian health data is digitized, with blockchain-secured electronic health records operational since 2007 [99]. The country stores genome data for 20 percent of its 1.3 million citizens, enabling precision medicine applications including AI-assisted cancer screening for breast, prostate, and colorectal cancers [99]. Estonia's digital health solutions have

cut operational healthcare costs by more than half, according to the Estonian Foresight Centre [99]. The nation is now exporting its model: Estonian firms collaborate with the NHS on data interoperability, AI radiology, and remote cardiac monitoring, showing that a small country's infrastructure choices can yield outsized influence on global healthcare AI adoption [102].

Risks and Limitations

While AI offers significant potential, responsible deployment requires careful attention to technical limitations, data quality issues, and systemic risks.

Cybersecurity: The Existential Threat

No risk more acutely threatens the Healthcare 2030 vision than the cybersecurity crisis facing interconnected health systems. On February 21, 2024, Change Healthcare (a clearinghouse processing approximately one-third of all U.S. healthcare transactions) suffered a ransomware attack by the ALPHV/BlackCat group that gained access via a Citrix portal lacking multifactor authentication [50]. By July 31, 2025, Change Healthcare had notified HHS that an estimated 192.7 million individuals were affected, approximately two-thirds of the U.S. population, making it the largest healthcare data breach ever recorded [48]. An AHA survey of approximately 1,000 hospitals found that 74 percent reported direct patient care impact, 94 percent reported financial impact, and 33 percent reported the attack disrupted more than half of their revenue [51].

Three months later, in May 2024, Ascension Health (one of the nation's largest nonprofit health systems with 140+ hospitals) was struck by the Black Basta ransomware group. The attack forced ambulance diversions, pharmacy closures, and a six-week loss of EHR access, affecting 5,599,699 individuals [52]. Patient volumes dropped 8–12 percent during the peak disruption period [50]. Healthcareransomware attacks increased approximately 45 percent in the first nine months of 2025 compared to the prior year, with an average breach cost of $10.22 million per incident (IBM, 2025) [53]. The average ransom demand targeting healthcare providers in 2025 was $514,000 [51].

AI systems introduce compounded security risks beyond traditional EHR vulnerabilities: model tampering, adversarial attacks on AI algorithms, data poisoning during training, and vastly expanded API-enabled attack surfaces created by FHIR data exchange and cloud-based AI platforms. The WHO has specifically identified cybersecurity risks to AI systems as vulnerabilities that "could endanger patient information or the trustworthiness of these algorithms and the provision of health care more broadly" [52]. Ransomware accounted for 69 percent of all stolen healthcare records in 2024, and the average time to identify and contain a healthcare breach was 291 days, the longest of any industry [51]. The connection between AI infrastructure and clinical operations means that a ransomware attack on an AI-dependent system does not just expose data. It can disable the clinical decision support, predictive monitoring, and automated alerting that patient safety increasingly depends upon.

The Equity Gap: Two Systems in One Future

The rural healthcare AI gap is stark and well-documented. A 2025 survey of 972 rural healthcare staff and leaders (Black Book Research) found that only 8 percent of rural community and critical access hospitals use AI-driven analytics for predictive healthcare, compared to 65 percent of U.S. hospitals overall [55]. Approximately half of rural hospitals currently operate with negative operating margins, leaving no financial cushion for technology investment [53]. Enterprise AI platforms range from $300,000 to $1 million or more; even basic AI tools carry costs of $40,000 or above [56]. Without a sustainable reimbursement pathway, the AI dividend will flow primarily to well-resourced urban health systems, precisely the inverse of where need is greatest.

Algorithmic bias represents a second, subtler equity threat. A landmark study by Obermeyer et al., reviewed in a 2025 *Nature* analysis, found that a widely used U.S. health risk prediction algorithm systematically underestimated the health needs of Black patients because historical cost data had embedded systemic barriers into the model. Black patients had 26.3 percent more chronic illnesses than white patients at the same risk score, yet were under-enrolled in care management programs. Recalibrating the algorithm nearly tripled Black patient enrollment [57]. A 2020 vaginal birth after cesarean (VBAC) calculator was found to predict systematically

lower success rates for Black and Hispanic women, encouraging unnecessary cesarean sections; the MFMU Network removed race from the algorithm in 2021 [58].

The Clinician Trust Gap and Pilot Purgatory

Despite 66 percent of physicians reporting AI use [23], only 16 percent use AI for clinical decision-making [24]. The Elsevier 2025 survey found only 32 percent of clinicians feel their institution provides adequate access to AI technologies, and only 30 percent have received sufficient training [24]. The Stanford-Harvard State of Clinical AI 2026 report, synthesizing 500+ medical AI studies, documented that clinicians following incorrect AI recommendations performed worse than if AI had not been used at all, confirming that over-reliance on flawed AI is a real and quantifiable patient safety risk [59].

“Pilot purgatory” is an institutional-level failure mode. The NHS case is instructive: 90 percent of NHS AI tools remain in pilot phases years after initial deployment, and 23 of 66 trusts that contracted for AI tools were still not using them in clinical practice as of June 2025 [45]. The friction isn’t primarily technical. It’s governance, procurement, workflow integration, and change management. A 2025 survey found that 96 percent of healthcare IT executives report facing significant challenges proving the ROI of AI investments, and 91 percent of small practices report hospital EHR systems are too complex for effective AI integration [54]. Organizations that have achieved scale (Kaiser Permanente, Mayo Clinic, Samsung Medical Center) did so by embedding AI governance into their operating model and maintaining longitudinal clinical partnerships over years and decades. Those that treat AI as an IT procurement exercise remain trapped in demonstration mode.

In every health system I’ve consulted with, the pattern is the same: the technology is ready before the organization is. Pilot purgatory isn’t a technology failure; it’s a leadership and governance failure, and it’s more common than any of us would like to admit.

Cost Barriers at the Point of Entry

The economics of AI adoption contain a structural barrier for smaller institutions. Enterprise AI implementations require not only the platform cost ($300,000–$1 million+) but EHR integration ($10,000–$50,000+ per system), HIPAA and regulatory validation (10–25 percent of project cost), and ongoing staff training [54]. A systemic review of 19 clinical AI cost-effectiveness studies published in *npj Digital Medicine* (2025) found that most clinical AI interventions are cost-effective or cost-saving in the long run [60]. Break-even is typically achievable within 12–24 months [60]. The up-front capital requirement effectively excludes many safety-net and rural hospitals from the market absent targeted policy intervention or grant support, such as the $50 billion Rural Health Transformation Program included in the July 2025 federal reconciliation bill [53].

Regulatory and Legal Landscape

The regulatory environment for AI-enabled medical devices continues to evolve as agencies worldwide develop frameworks that balance innovation with patient safety.

FDA: Governance Infrastructure Expanding

The FDA's regulatory framework for AI-enabled medical devices has undergone its most significant evolution since the agency first authorized an AI device in the 1990s. The December 2024 finalization of the Predetermined Change Control Plan (PCCP) framework enables AI device manufacturers to prespecify algorithm modifications that can be implemented without new regulatory submissions, provided they execute precisely as planned [22]. This addresses the fundamental tension between AI's iterative, continuously improving nature and a regulatory system historically designed for static devices. The PCCP framework applies across the 510(k), De Novo, and PMA pathways and is being developed in coordination with Health Canada and the UK MHRA (Medicines and Healthcare products Regulatory Agency) [22].

The FDA's post-market monitoring infrastructure has not kept pace with authorization volume. Approximately 6 percent of AI-enabled devices have

been recalled, often within the first year, pointing to limitations in pre-market clinical testing [20]. FDA staffing was down approximately 2,500 employees (approximately 15 percent) from 2023 levels as of September 2025, constraining real-world monitoring capacity [20]. In September 2025, the FDA issued a formal Request for Public Comment on best practices for measuring and evaluating real-world AI device performance, including managing performance drift, an acknowledgment that the monitoring infrastructure remains a work in progress [22].

EU AI Act: High-Risk Classification and Compliance Burden

The European Union's AI Act, which entered into force on August 1, 2024, classifies all medical AI devices regulated under the EU MDR and IVDR as high-risk AI systems [61]. Full compliance for new high-risk AI market placements is required by August 2, 2026, with AI integrated into regulated products (including medical devices) required to comply by August 2, 2027 [61]. Compliance requirements include thorough technical documentation, strong risk management specific to AI risks, data governance procedures covering bias detection and data representativeness, mandatory post-market surveillance, and registration in the EU AI database [59]. Noncompliance carries penalties of up to 7 percent of annual worldwide turnover (or €35 million) for prohibited practices, and up to 3 percent of annual worldwide turnover [59].

The compliance cost burden is specific and documented. *JMIR AI* (October 2025) estimates ongoing EU AI Act compliance costs at approximately €29,277 annually per AI unit, with certification costs of €16,800–€23,000 per unit [62]. These costs are layered atop existing EU MDR compliance requirements. Notified body capacity (already a bottleneck under MDR implementation) faces further strain as all Class IIa and above medical AI devices must now address both sets of requirements simultaneously [61]. That's a meaningful burden, and smaller vendors aren't going to absorb it easily.

Reimbursement: The Remaining Bottleneck

As of early 2026, only 26 CPT codes cover AI services, with just 3 at Category I level [63]. This reimbursement gap is the single most cited

barrier to AI adoption in primary-care settings [64]. In January 2025, CMS introduced new billing codes for FDA-cleared digital mental health treatments, and the September 2025 finalization of the Hospital Outpatient Prospective Payment System for 2026 introduced new payment mechanisms for remote patient monitoring and AI-assisted radiology workflows [65]. The Health Tech Investment Act, introduced in Congress in 2025, would establish a dedicated pathway for AI-specific coverage determinations. CMS's ACCESS Model, introduced in 2025, creates performance-based incentives for qualifying health systems to adopt AI tools demonstrating measurable clinical improvements [63].

These legislative and coding developments signal that the reimbursement infrastructure is evolving. The pace of change has not matched the pace of clinical deployment.

Interoperability: TEFCA and FHIR Progress

The interoperability infrastructure that AI systems require to function is materially improving. As of early 2026, more than 71,000 sites or organizations participate in Trusted Exchange Framework and Common Agreement (TEFCA) through 11 qualified health information networks [66]. Fast Healthcare Interoperability Resources (FHIR) is now actively used in 71 percent of countries for at least some use cases, with 73 percent of countries that have electronic health data exchange regulations mandating or advising FHIR use [67]. ONC enforcement of information blocking rules has escalated, with HHS OIG imposing civil monetary penalties of up to $1 million per violation [68].

Interoperability barriers persist. Thirty-five percent of physicians still rely on fax or e-fax to exchange protected health information [69]. Many organizations operate 20 or more different software systems simultaneously, with approximately 20 percent of patients potentially having incorrect medical records due to poor integration [67]. True semantic interoperability (where codes, units, and terminologies align across organizations) remains an aspirational goal in many healthcare markets, limiting the quality of data available for AI training and inference.

Malpractice and Liability

Malpractice claims involving AI tools rose 14 percent between 2022 and 2024, with most cases originating in radiology, oncology, and cardiology, the specialties with the highest concentration of FDA-cleared AI diagnostic tools [70]. The legal standard is still forming: courts currently apply a "reasonable physician" standard augmented with criteria including documented AI training, adherence to clinical guidelines, exercise of independent judgment, on-label AI use, and quality of documentation explaining acceptance or override of AI recommendations [68]. Some malpractice insurers now include AI-specific exclusions or require AI training as a coverage condition [68]. The AMA survey found that 87 percent of physicians rate "not being held liable for AI model errors" as a critical prerequisite for adoption [23].

Ethical Implications for Health Equity

Effective AI deployment must address fundamental questions of fairness, transparency, and equitable access to ensure that technology benefits all populations.

The AI Divide Risk

The most consequential ethical question facing Healthcare 2030 is not whether AI works. The evidence base is clear that it does, across a growing range of applications. The question is whether AI will narrow or widen the health disparities that already define American and global healthcare. That's the 800-pound gorilla in the room, and we haven't resolved it. The risk of a two-tier system is real: well-resourced academic health systems capturing the efficiency and quality dividends of AI while safety-net hospitals, rural facilities, and underserved communities are left behind. As Rebecca Mishuris, MD, MPH, chief medical information officer at Mass General Brigham, stated: "Many healthcare systems across the country cannot afford to pilot or implement generative AI, which leaves out providers and patients from potential benefits like reduced provider burnout and improved patient experiences" [69].

The Obermeyer algorithmic bias study and the VBAC calculator recalibration are documented cases where AI systems encoded existing disparities into clinical decision-making [57,58]. A 2025 systematic review found that AI health algorithms "present serious challenges to patient privacy, data security, institutional transparency, and health equity" and carry the potential to perpetuate health inequities if not actively designed against [72]. The WHO has identified inclusive data practices, equitable governance, participatory design, and local capacity-building as the four interdependent domains that determine whether AI bridges or widens disparities [73].

The counter-evidence is also significant. Autonomous FDA-cleared AI tools for diabetic retinopathy screening (LumineticsCore) are already delivering specialist-level screening in rural primary-care settings where ophthalmologists are unavailable [74]. The Penda Health/OpenAI AI clinical copilot study in Nairobi showed a 16 percent relative reduction in diagnostic errors and a 13 percent reduction in treatment errors across 39,849 patient visits, in a lower-resource setting where AI's equity dividend is most needed [75].

Data Governance and Patient Ownership

The governance of the health data that AI systems require is among the most complex unresolved questions of this era. We don't yet have a coherent national framework in the United States. The European Health Data Space (EHDS), which entered into force on March 26, 2025, establishes the most thorough harmonized framework for health data access and patient rights in the world, including the right to access, share, restrict, audit, and correct one's health data, and an explicit opt-out from secondary data use [76]. U.S. law does not provide patients with property rights over their health data, though HIPAA confers access rights.

Nearly half of healthcare organizations that use generative AI have no formal approval process for AI adoption, and only 31 percent actively monitor deployed AI systems [77]. California's Assembly Bill 3030 requires healthcare providers to disclose AI use in patient care and obtain explicit consent before deploying AI-powered systems, and more than half of U.S. states have introduced or passed legislation specifically addressing healthcare AI [75]. The patchwork of state-level regulation, absent a

coherent federal framework, creates compliance complexity that may itself slow beneficial AI deployment.

Algorithmic Accountability

When an AI system makes a clinical error, the question of who bears responsibility remains legally and ethically unresolved. The physician who accepted the AI's recommendation? The hospital that deployed an inadequately validated tool? The software vendor who designed a biased training dataset? The emerging legal standard (documented AI training, independent judgment, on-label use, thorough documentation) places primary responsibility on the clinical user, even when the underlying error is the algorithm's [68]. This structure creates a liability environment that may paradoxically discourage the most conscientious clinicians from deploying AI tools, precisely because they are more likely to document their reasoning and therefore expose their decision-making to scrutiny.

The American College of Physicians has called for AI to "enhance human intelligence, not supplant it," recommending that AI be limited to a supportive role in clinical decision-making, with accountability remaining with the physician [78]. The Joint Commission and the Coalition for Health AI (CHAI), in their September 2025 joint guidance (the first full U.S. framework for responsible AI adoption in health systems), identified ethical evaluation, usefulness, effectiveness, and financial performance as four core governance pillars [75].

Preserving the Therapeutic Relationship

The concern most consistently voiced by physicians across major surveys, the potential loss of the human element in patient care, deserves serious engagement rather than dismissal [79]. The evidence from ambient AI scribe deployment is instructive: at the Permanente Medical Group, 47 percent of patients reported their doctor spent less time on the computer, 39 percent said their doctor spent more time speaking directly with them, and 56 percent reported a positive impact on visit quality, with no patients reporting a negative impact [40]. The mechanism is precisely the opposite of what critics fear: AI eliminating documentation burden frees clinical attention for the relational dimension of medicine that AI can't replicate.

Medical education is adapting. The AAMC reported that medical schools incorporating AI into their curricula increased from 53 percent in 2023 to 77 percent in 2024, with a national AI competency development project launched in July 2025 to standardize training across the learning continuum [80]. Starting fall 2025, Stanford Medicine required all medical and physician assistant students to complete AI literacy training covering clinical application, ethical implications, and critical evaluation of AI outputs [81]. Mayo Clinic Medical School introduced a "Human Skills in the Age of AI" curriculum that explicitly pairs AI literacy with the communication, empathy, and ethical decision-making skills that constitute the irreducible core of clinical professionalism [78].

Leadership Implementation Framework

Healthcare leaders preparing for AI integration should adopt a structured approach that aligns technology capabilities with organizational strategy and clinical priorities.

The Strategic Entry Point: AI Documentation

For health system leaders navigating the path to Healthcare 2030, the evidence strongly supports a specific sequencing strategy. The entry point that delivers the fastest return, builds clinician trust, and creates the cultural foundation for subsequent AI deployment is ambient clinical documentation. The ROI is rapid: clinical documentation AI achieves break-even within three to six months, saving two to three hours per physician per day and generating $50,000–$150,000 in annual value per physician through recovered productivity [54]. The multicenter burnout study found that 30 days of ambient AI scribe use reduced burnout from 51.9 percent to 38.8 percent [19]. Northwestern Medicine physicians using DAX Copilot in at least 50 percent of encounters saw 11.3 additional patients per month, 24 percent less time in notes, and a 17 percent decrease in after-hours documentation [82].

The University of Wisconsin Health randomized trial, published in *NEJM AI*, provides the strongest evidence to date: a 30-minute reduction in documentation time per provider per day, clinically meaningful burnout score reductions, improved note billing accuracy, and improved measures of

well-being [18]. This is not a pilot finding. It's evidence that should inform every health system's capital allocation decision.

The Value-Based Care Transition

CMS's goal of 100 percent of Traditional Medicare beneficiaries in accountable-care relationships by 2030 creates both a deadline and an infrastructure requirement for AI deployment [27]. Beginning January 1, 2026, the Transforming Episode Accountability Model (TEAM) requires hospitals to manage total care costs for specific procedures, a structural incentive that rewards exactly the capabilities AI enables: predicting high-risk patients, managing transitions, reducing readmissions, and closing care gaps [83]. The Medicare Shared Savings Program achieved its highest savings to date in 2024, marking seven consecutive years of cost reductions and confirming that value-based models produce measurable results [81].

AI tools are the operational enabler of the accountable care transition. Kaiser Permanente's AAM shows that predictive analytics preventing 500+ deaths per year is not a research result. It's a documented operational outcome [41]. CommonSpirit Health's AI-powered care gap closure resulted in 61,000+ care gap orders in FY2025, a five-fold increase over the prior year [84]. Nebraska Medicine's AI-driven capacity management achieved a 5 percent reduction in average length of stay, generating the equivalent of 37 additional inpatient beds without capital construction [82].

Hospital-at-Home as the Care Delivery Model

As of October 2024, 366 hospitals participated in the CMS Acute Hospital Care at Home initiative across 38 states, serving more than 31,000 patients [85]. The September 2024 CMS study documented lower 30-day mortality rates, lower hospital-acquired condition rates across all six types evaluated, and approximately 20 percent less Medicare spending in the 30-day post-discharge period compared to traditional inpatient care [83]. Congress has extended the program through the end of 2030 [86]. Advocate Health has freed up 33,000 hospital bed days by caring for more than 9,400 patients at home since 2020 [87].

AI-powered remote patient monitoring is the clinical infrastructure that makes hospital-at-home safe and scalable. A deep learning model using

chest-worn wearable data from 888 non-ICU inpatients, published in *Nature Communications* (November 2025), predicted in-hospital deterioration with a 24-hour lead time and an AUC of 0.89, detecting 126 more alerts than manual EHR monitoring with a 105-minute average lead-time advantage [88].

Building the AI Workforce

The shift to Healthcare 2030 requires workforce development at every level. In my experience, the biggest gap isn't technology. It's the human infrastructure needed to deploy it well. The AAMC's national AI competency project, launched in July 2025 with Josiah Macy Jr. Foundation support, is developing standardized AI training across undergraduate, graduate, and continuing medical education, a landmark shift in how the profession prepares its members [78]. New role categories are emerging rapidly: AI-supported radiology assistants, clinical decision support analysts, healthcare AI prompt engineers, clinical informaticists specializing in AI, and AI ethics and compliance specialists [89]. As of early 2026, `Indeed.com` lists more than 1,314 "Clinical AI Specialist" job postings [87].

Phased Implementation Pathway

The evidence supports a four-phase implementation sequence aligned with institutional readiness and risk tolerance. Phase 1 (ambient AI documentation) addresses burnout, demonstrates ROI rapidly, and builds clinician trust. Phase 2 (clinical decision support) introduces AI-assisted diagnosis and treatment recommendations in structured, human-supervised contexts. Phase 3 (predictive analytics and population health management) deploys AI for high-risk patient identification, care gap closure, and epidemiological monitoring. Phase 4 (autonomous systems) applies AI to tasks where human judgment is supplemented by AI action: robotic surgery assistance, continuous remote monitoring with automated alerts, AI-managed triage. This phased approach mirrors the validated maturity pathway demonstrated by the institutions studied in this chapter.

Future Outlook

The next five years will be critical in determining whether AI becomes a standard tool for improving healthcare delivery or remains concentrated in advanced academic centers.

Ambient Clinical Intelligence as the New Standard

By 2030, ambient clinical intelligence (the real-time AI capture, synthesis, and integration of clinical encounters) will be the expected baseline of clinical practice, not an innovation. It won't be optional any more than the EHR is optional today. Nuance DAX Copilot was adopted by more than 400 organizations within its first year of commercial availability and is now deployed across expansion markets in Austria, France, Germany, and Ireland [90]. Abridge has been deployed to 24,000 Kaiser Permanente physicians and is expanding across dozens of additional major health systems [39]. More than $1 billion was invested in the AI scribe industry in 2025 alone [19]. The trajectory points to a 2030 in which ambient AI documentation is as universal as the EHR itself.

AI Drug Discovery: Compressing the Decade

The shift in pharmaceutical R&D is proceeding on a trajectory that'll produce the first fully AI-designed drugs in phase 3 trials by 2030. Traditional drug discovery takes approximately 15 years from target identification to approval at an average cost of approximately $1–2 billion per drug [91]. AI-native firms are reporting Phase I success rates of 80-90 percent, versus the historical industry average of approximately 50 percent [89]. Insilico Medicine advanced its IPF candidate ISM001-055 from target identification to a preclinical candidate in approximately 18 months at a cost of approximately $150,000, a process that traditionally takes four to six years and tens of millions of dollars [92]. AlphaFold, the Google DeepMind protein structure prediction system that won the 2024 Nobel Prize in Chemistry, has predicted the structure of more than 200 million proteins and is being used by more than 2 million researchers globally [93].

Digital Twins: From Concept to Clinical Practice

Digital twins in healthcare (virtual models integrating genomics, imaging, wearable data, and clinical records for real-time patient simulation) are advancing from research concept to clinical application. The healthcare digital twin market, valued at approximately $2.1 billion in 2024, is projected to grow at approximately 27 percent annually [94]. The U.S. digital twin healthcare market is expected to reach $27.34 billion by 2035 (U.S. segment) at a CAGR of 41.34 percent [92]. Johns Hopkins University pioneered the first FDA-approved virtual heart model, enabling simulation-based treatment customization [95]. Mater Hospital Dublin's operational digital twin cut CT/MRI patient wait times by approximately 240 minutes and increased scan capacity by 26–32 percent [93]. A Roche expert at the 2025 Outsourcing in Clinical Trials DACH conference projected that fully individualized patient digital twins for clinical decision-making are "five to seven years away," placing solid tools squarely within the 2030 planning horizon [92].

Foundation Models and Medical Intelligence

Epic Systems and Microsoft Research unveiled CoMET in 2025, a family of massive generative models trained on more than 300 million patient records and 16 billion encounters, capable of simulating future health timelines to predict medical events [96]. Google's AMIE research system, evaluated in a blinded OSCE study, showed greater diagnostic accuracy than primary-care physicians on 28 of 32 evaluation axes and achieved a top-10 diagnostic accuracy of 59.1 percent versus 33.6 percent for unassisted PCPs [97]. The Penda Health/OpenAI real-world clinical copilot study documented a 16 percent relative reduction in diagnostic errors in 39,849 patient visits [73]. Microsoft's MAI-DxO platform, using a "chain-of-debate" multi-agent architecture, has reported diagnostic accuracy exceeding 85 percent in research case studies [98].

Gartner projects that AI agents and virtual agents could manage up to 75 percent of routine healthcare interactions by 2030, a projection that reflects the convergence of natural language capability, EHR integration, and the sheer volume of administrative tasks that consume current clinical capacity [99]. The primary care panel of the future, as the Lumeris 2030 vision articulates, may expand from 2,000 to 5,000+ patients per physician, with

AI handling approximately 80 percent of care coordination, medication refills, scheduling, and basic health coaching. Not replacing the physician, but multiplying their reach [100].

Key Takeaways

The $187.7 billion market is a consequence, not a cause. The global AI in healthcare market's projected growth to $187.7 billion by 2030 [1] reflects genuine clinical value being delivered at scale. The investments being made now are responding to documented workforce shortages, unsustainable cost trajectories, and proven ROI across multiple application domains.

The workforce crisis is the most powerful accelerator. A projected U.S. physician shortage of up to 86,000 by 2036 [3] and a global health workforce deficit of approximately 10 million workers by 2030 [4] mean that AI is not competing with clinicians. It's the mechanism by which healthcare systems will be able to serve an aging population with a constrained workforce.

Ambient AI documentation is the highest-confidence entry point. Across multiple peer-reviewed studies, ambient clinical documentation AI consistently shows reduced burnout (51.9 percent to 38.8 percent) [19], time savings of 30 minutes per day [18], improved patient satisfaction [40], and break-even within three to six months [54]. No other AI application offers this convergence of clinical evidence, physician acceptance, and financial return.

Cybersecurity is a prerequisite, not an afterthought. The Change Healthcare breach affecting 192.7 million individuals [48] and healthcare ransomware increasing 30 percent in 2025 [51] show that AI investment without commensurate cybersecurity investment is building on an unstable foundation. AI infrastructure widens the attack surface and must be secured accordingly.

The equity gap is the defining ethical challenge of healthcare 2030. With only 8 percent of rural hospitals using AI analytics versus 65 percent overall [53] and algorithmic bias documented to systematically disadvantage Black and Hispanic patients [55,56], the distribution of

AI's benefits is as important as the development of AI's capabilities. Healthcare 2030 can be the most equitable or the most unequal period in American healthcare history. The outcome is determined by deliberate policy choices being made now.

Reimbursement is the critical bottleneck. Only 26 CPT codes cover AI services, with just 3 at Category I level [61]. Until a sustainable reimbursement pathway exists, AI adoption will remain concentrated in well-resourced institutions that can afford multiyear ROI horizons. The Health Tech Investment Act and CMS's ACCESS model represent the most promising legislative pathways to closing this gap.

The regulatory infrastructure is maturing but incomplete. The FDA's PCCP framework [22], the EU AI Act's high-risk classification with full compliance timelines through 2027 [59], the WHO's 40+ recommendations for AI governance [54], and IMDRF's Good Machine Learning Practice guidelines [22] collectively represent unprecedented international regulatory coordination. Post-market surveillance remains the weakest link.

Scaling is harder than building. The NHS's 90 percent pilot rate [45] and the documented gap between research performance and real-world clinical deployment [57] confirm that governance, integration, workflow redesign, and change management determine whether AI reaches patients or remains confined to demonstrations. The institutions that have achieved scale (Kaiser Permanente, Mayo Clinic, Samsung Medical Center) treat AI implementation as an organizational capability, not an IT project.

The "augmentation" model is evidence-based; the "replacement" debate is settled. The American College of Physicians' formal position, the Stanford-Harvard synthesis of 500+ studies [57], and the real-world results from AMIE, Penda Health, and ambient scribes converge on the same conclusion: AI delivers greatest value as a "centaur," augmenting human clinical judgment rather than operating autonomously. The therapeutic relationship, when AI is deployed correctly, is enhanced, not diminished.

2030 is a planning deadline, not a prediction. The institutions already operating at the Healthcare 2030 level (Samsung Medical

Center, Kaiser Permanente, Humber River Health System, Estonia's national health system) did not arrive there by reacting to trends. They invested in infrastructure, governance, and clinical-AI partnerships over years and decades. Healthcare leaders who begin building that infrastructure today will be positioned to deliver the vision. Those who wait for certainty will arrive too late to capture the value and too unprepared to manage the risks.

References

1. Grand View Research. AI in healthcare market size to reach $187.7 billion by 2030 at CAGR 38.5%. PR Newswire; 2025 Apr 28 [cited 2026 Jul 4]. Available from: https://www.prnewswire.com/news-releases/ai-in-healthcare-market-size-to-reach-187-7-billion-by-2030-at-cagr-38-5---grand-view-research-inc-302439558.html.

2. Centers for Medicare & Medicaid Services, Office of the Actuary. CMS Office of the Actuary releases 2021-2030 projections of national health expenditures. Baltimore (MD): CMS; 2022 Mar 28 [cited 2026 Jul 4]. Available from: https://www.cms.gov/newsroom/press-releases/cms-office-actuary-releases-2021-2030-projections-national-health-expenditures.

3. Association of American Medical Colleges. New AAMC report shows continuing projected physician shortage. Washington (DC): AAMC; 2024 Mar 21 [cited 2026 Jul 4]. Available from: https://www.aamc.org/news/press-releases/new-aamc-report-shows-continuing-projected-physician-shortage.

4. Boniol M, Kunjumen T, Nair TS, Siyam A, Campbell J, Diallo K. The global health workforce stock and distribution in 2020 and 2030: a threat to equity and 'universal' health coverage? BMJ Glob Health. 2022; 7(6):e009316. doi:10.1136/bmjgh-2022-009316.

5. Kaiser Permanente. Artificial intelligence. Oakland (CA): Kaiser Permanente; 2025 [cited 2026 Jul 4]. Available from:

https://about.kaiserpermanente.org/expertise-and-impact/public-policy/our-key-issues/artificial-intelligence.

6. Insight Partners. How the Mayo Clinic is transforming healthcare with AI. 2025 [cited 2026 Jul 4]. Available from: https://www.insightpartners.com/ideas/how-the-mayo-clinic-is-transforming-healthcare-with-ai.

7. Fairbairn TA, Mullen L, Nicol E, Lip GYH, Schmitt M, Shaw M, et al. Implementation of a national AI technology program on cardiovascular outcomes and the health system. Nat Med. 2025; 31:1136-45. doi:10.1038/s41591-025-03620-y.

8. Healthcare Information and Management Systems Society. Samsung Medical Center earns 2024 Davies Award, first ever in Republic of Korea. Chicago (IL): HIMSS; 2024 [cited 2026 Jul 4]. Available from: https://www.himss.org/news-center/samsung-medical-center-earns-2024-davies-award-first-ever-republic-korea.

9. Parikh H, Nighan M, Montague V, Mutanuka-Hiebel M, Yang A. The Change Healthcare cybersecurity breach. Nixon Peabody LLP; 2025 Nov 12 [cited 2026 Jul 4]. Available from: https://www.nixonpeabody.com/insights/alerts/2025/11/12/change-healthcare-cybersecurity-breach-impact-on-healthcare-providers.

10. Elsevier. Elsevier survey reveals growing clinician AI adoption despite trust and training gaps. HLTH; 2025 Jul 16 [cited 2026 Jul 4]. Available from: https://hlth.com/insights/news/elsevier-survey-reveals-growing-clinician-ai-adoption-despite-trust-and-training-gaps-2025-07-16.

11. MarketsandMarkets. US artificial intelligence (AI) in healthcare market worth $43.30 billion by 2030. 2025 [cited 2026 Jul 4]. Available from: https://www.marketsandmarkets.com/PressReleases/us-artificial-intelligence-healthcare.asp.

12. KPMG. Generative AI is poised to transform healthcare. 2024 [cited 2026 Jul 4]. Available from: https://assets.kpmg.com/content/dam/kpmgsites/sg/pdf/2024/08/

generative-ai-poised-transform-healthcare.pdf.coredownload.inline.pdf.

13. McKinsey & Company. The coming evolution of healthcare AI: toward a modular architecture. 2025 [cited 2026 Jul 4]. Available from: https://www.mckinsey.com/industries/healthcare/our-insights/the-coming-evolution-of-healthcare-ai-toward-a-modular-architecture.

14. McKinsey Global Institute. AI: work partnerships between people, agents, and robots. 2025 [cited 2026 Jul 4]. Available from: https://www.mckinsey.com/mgi/our-research/agents-robots-and-us-skill-partnerships-in-the-age-of-ai.

15. Association of American Medical Colleges. New research shows increasing physician shortages in both primary and specialty care. Washington (DC): AAMC; 2018 [cited 2026 Jul 4]. Available from: https://www.aamc.org/news/press-releases/new-research-shows-increasing-physician-shortages-both-primary-and-specialty-care.

16. Bain & Company. Primary care 2030. 2022 [cited 2026 Jul 4]. Available from: https://www.bain.com/insights/primary-care-2030.

17. TNAA. The state of healthcare workforce shortages: challenges, causes, and solutions. 2025 [cited 2026 Jul 4]. Available from: https://tnaa.com/blog/the-state-of-healthcare-workforce-shortages-challenges-causes-and-solutions.

18. University of Wisconsin School of Medicine and Public Health. Studies find AI technology for clinical documentation aids efficiency and reduces burnout. Madison (WI): UW; 2025 [cited 2026 Jul 4]. Available from: https://www.med.wisc.edu/news/ambient-ai-improves-practitioner-well-being.

19. O'Reilly KB. How much can ambient AI scribes help cut doctor burnout? Reporting on a JAMA Network Open multicenter study. American Medical Association; 2025 Oct 2 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/physician-

health/how-much-can-ambient-ai-scribes-help-cut-doctor-burnout.

20. Bipartisan Policy Center. FDA oversight: understanding the regulation of health AI tools. Washington (DC): BPC; 2025 [cited 2026 Jul 4]. Available from: https://bipartisanpolicy.org/issue-brief/fda-oversight-understanding-the-regulation-of-health-ai-tools.

21. Almarie B, Gonzalez-Gonzalez LF, dos Santos Barbosa LA, Lutz A, Grosse U, Fregni F. Machine learning-enabled medical devices authorized by the US Food and Drug Administration in 2024: regulatory characteristics, predicate lineage, and transparency reporting. Biomedicines. 2025; 13(12):3005. doi:10.3390/biomedicines13123005.

22. McDermott+. FDA issues final guidance on predetermined change control plans for AI-enabled devices. 2024 [cited 2026 Jul 4]. Available from: https://www.mcdermottplus.com/insights/fda-issues-final-guidance-on-predetermined-change-control-plans-for-ai-enabled-devices.

23. Henry TA. 2 in 3 physicians are using health AI—up 78% from 2023. American Medical Association; 2025 Feb 26 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/2-3-physicians-are-using-health-ai-78-2023.

24. Elsevier. Clinician of the future 2025 [survey of 2,206 clinicians across 109 countries]. See reference 10. HLTH; 2025 [cited 2026 Jul 4]. Available from: https://hlth.com/insights/news/elsevier-survey-reveals-growing-clinician-ai-adoption-despite-trust-and-training-gaps-2025-07-16.

25. athenahealth, athenaInstitute. AI on the frontlines of care: insights from clinicians. 2025 [cited 2026 Jul 4]. Available from: https://www.athenahealth.com/athenainstitute/perspectives/how-clinicians-feel-ai-adoption-in-healthcare-frontlines-of-care.

26. Zimlichman E, Nicklin W, Aggarwal R, Bates DW. Health care 2030: the coming transformation. NEJM Catal Innov Care Deliv. 2021.

doi:10.1056/CAT.20.0569. Available from: https://foh.health/wp-content/uploads/2024/09/Healthcare-2030_NEJM-Catalyst_2021.pdf [cited 2026 Jul 4].

27. Centers for Medicare & Medicaid Services. CMS Innovation Center's strategy to support high-quality primary care. Baltimore (MD): CMS; 2023 [cited 2026 Jul 4]. Available from: https://www.cms.gov/blog/cms-innovation-centers-strategy-support-high-quality-primary-care.

28. Deloitte. Life sciences and health care predictions 2030. 2025 [cited 2026 Jul 4]. Available from: https://www.deloitte.com/uk/en/Industries/life-sciences-health-care/collections/life-sciences-and-health-care-predictions.html.

29. PwC. From breaking point to breakthrough: the $1 trillion opportunity to reinvent healthcare. 2025 [cited 2026 Jul 4]. Available from: https://www.pwc.com/us/en/industries/health-industries/library/future-of-health.html.

30. Boston Consulting Group. How digital and AI will reshape health care in 2025. 2025 [cited 2026 Jul 4]. Available from: https://www.bcg.com/publications/2025/digital-ai-solutions-reshape-health-care-2025.

31. Kriwet C. 3 ways AI will change healthcare by 2030. World Economic Forum; 2020 Jan [cited 2026 Jul 4]. Available from: https://www.weforum.org/stories/2020/01/future-of-artificial-intelligence-healthcare-delivery.

32. CB Insights. The future according to Mayo Clinic: how AI is transforming the hospital. Mayo Clinic Business Development; 2024 [cited 2026 Jul 4]. Available from: https://businessdevelopment.mayoclinic.org/wp-content/uploads/2024/11/The-future-according-to-Mayo-Clinic-How-AI-is-transforming-the-hospital-CB-Insights.pdf.

33. Mayo Clinic. Researchers develop AI-ECG to detect obstructive sleep apnea [ECG-OSA study, JACC: Advances]. Rochester (MN): Mayo

Clinic; 2026 [cited 2026 Jul 4]. Available from: https://www.mayoclinic.org/medical-professionals/cardiovascular-diseases/news/mayo-clinic-cardiovascular-medicine-artificial-intelligence-groups-researchers-develop-ai-ecg-to-detect-obstructive-sleep-apnea/mac-20594007.

34. Mayo Clinic. AI-based ACE score accurately predicts decompensation and liver outcomes in patients diagnosed with cirrhosis [JHEP Reports]. Rochester (MN): Mayo Clinic; 2025 [cited 2026 Jul 4]. Available from: https://www.mayoclinic.org/medical-professionals/digestive-diseases/news/ai-based-ace-score-accurately-predicts-decompensation-and-liver-outcomes-in-patients-diagnosed-with-cirrhosis/mac-20593551.

35. Yu Y, Hu X, Rajaganapathy S, Feng J, Abdelhameed A, Li X, et al. Accelerating AI innovation in healthcare: real-world clinical research applications on the Mayo Clinic Platform. npj Health Syst. 2026. doi:10.1038/s44401-026-00068-1. PMID: 41710561; PMCID: PMC12909107. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12909107.

36. Mayo Clinic launches AI-powered clinical development platform using international partner ecosystem [Mayo Clinic Platform_Orchestrate]. Interhospi; 2025 [cited 2026 Jul 4]. Available from: https://interhospi.com/mayo-clinic-launches-ai-powered-clinical-development-platform-using-international-partner-ecosystem.

37. Cleveland Clinic, Palantir forge AI pact to improve patient flow. Healthcare Dive; 2022 [cited 2026 Jul 4]. Available from: https://www.healthcaredive.com/news/Cleveland-Clinic-Palantir-AI-patient-flow/637937.

38. Mihaljevic T. 5 AI tools making an impact today [Cleveland Clinic AI scribe: 2 million patient encounters]. LinkedIn; 2026 [cited 2026 Jul 4]. Available from: https://www.linkedin.com/posts/tmihaljevic_5-ai-tools-making-an-impact-today-activity-7430234833215348736-wWGE.

39. Cleveland Clinic, IBM, and the Hartree Centre collaborate to advance healthcare and life sciences through artificial intelligence and quantum computing. Cleveland Clinic Newsroom; 2024 Jun 6 [cited 2026 Jul 4]. Available from: https://newsroom.clevelandclinic.org/2024/06/06/cleveland-clinic-ibm-and-the-hartree-centre-collaborate-to-advance-healthcare-and-life-sciences-through-artificial-intelligence-and-quantum-computing.

40. Cleveland Clinic accelerates clinical trial recruitment with roll-out of Dyania Health's artificial intelligence platform across health system. Cleveland Clinic Newsroom; 2025 Aug 27 [cited 2026 Jul 4]. Available from: https://newsroom.clevelandclinic.org/2025/08/27/cleveland-clinic-accelerates-clinical-trial-recruitment-with-roll-out-of-dyania-healths-artificial-intelligence-platform-across-health-system.

41. The Permanente Federation. Quality assurance informs large-scale use of ambient AI clinical documentation. 2025 [cited 2026 Jul 4]. Available from: https://permanente.org/quality-assurance-informs-large-scale-use-of-ambient-ai-clinical-documentation.

42. American Medical Association. AI scribes save 15,000 hours, and restore the human side of medicine. 2025 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/ai-scribes-save-15000-hours-and-restore-human-side-medicine.

43. American Medical Association. Kaiser Permanente's AI approach puts patients and doctors first [interview with Vincent Liu]. 2023 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/kaiser-permanente-s-ai-approach-puts-patients-and-doctors-first.

44. American Medical Association. Kaiser Permanente's AI approach puts patients and doctors first [NLP inbox management; same source as reference 43]. 2023 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-

health/kaiser-permanente-s-ai-approach-puts-patients-and-doctors-first.

45. BHF Data Science Centre. Groundbreaking AI trained on de-identified patient data to predict healthcare needs. 2025 [cited 2026 Jul 4]. Available from: https://bhfdatasciencecentre.org/news-and-events/groundbreaking-ai-trained-on-de-identified-patient-data-to-predict-healthcare-needs.

46. Department of Health and Social Care (UK). AI to be trialled at unprecedented scale across NHS screening. London: GOV.UK; 2025 Sep 22 [cited 2026 Jul 4]. Available from: https://www.gov.uk/government/news/ai-to-be-trialled-at-unprecedented-scale-across-nhs-screening.

47. Nelson Advisors. AI in the NHS: transformative promise v's operational reality. 2025 [cited 2026 Jul 4]. Available from: https://www.healthcare.digital/single-post/ai-in-the-nhs-transformative-promise-v-s-operational-reality.

48. HealthTechX Asia. Samsung Medical Centre's path to smart healthcare. 2024 [cited 2026 Jul 4]. Available from: https://www.healthtechx-asia.com/articles/samsung-medical-centres-path-smart-healthcare.

49. Southwick R. Samsung Medical Center offers a glimpse of the future of health care (HIMSS 2025). Chief Healthcare Executive; 2025 [cited 2026 Jul 4]. Available from: https://www.chiefhealthcareexecutive.com/view/samsung-medical-center-offers-a-glimpse-of-the-future-of-health-care-himss-2025.

50. Nixon Peabody LLP. The Change Healthcare cybersecurity breach [same source as reference 9]. 2025 [cited 2026 Jul 4]. Available from: https://www.nixonpeabody.com/insights/alerts/2025/11/12/change-healthcare-cybersecurity-breach-impact-on-healthcare-providers.

51. American Hospital Association. Change Healthcare cyberattack underscores urgent need to strengthen cyber preparedness [survey of

nearly 1,000 hospitals]. Washington (DC): AHA; 2024 Mar [cited 2026 Jul 4]. Available from: https://www.aha.org/system/files/media/file/2025/02/Change-Healthcare-Cyberattack-Underscores-Urgent-Need-to-Strengthen-Cyber-Preparedness.pdf.

52. Alder S. Ascension ransomware attack affects 5.6 million patients. HIPAA Journal; 2024 Dec 20 [cited 2026 Jul 4]. Available from: https://www.hipaajournal.com/ascension-cyberattack-2024.

53. Neagu V. Hospitals under attack: ransomware in healthcare in 2025. CyberGlobal; 2025 Nov 4 [cited 2026 Jul 4]. Available from: https://cybergl.com/blog/ransomware-in-healthcare.

54. World Health Organization. WHO releases AI ethics and governance guidance for large multi-modal models. Geneva: WHO; 2024 Jan 18 [cited 2026 Jul 4]. Available from: https://www.who.int/news/item/18-01-2024-who-releases-ai-ethics-and-governance-guidance-for-large-multi-modal-models.

55. Healthcare Brew. Rural healthcare crosses fingers for AI investment, reimbursement. 2025 Sep 17 [cited 2026 Jul 4]. Available from: https://www.healthcare-brew.com/stories/2025/09/17/rural-healthcare-ai-investment-reimbursement.

56. Murphi.ai. Cost of implementing AI in healthcare 2025. 2025 [cited 2026 Jul 4]. Available from: https://murphi.ai/cost-of-implementing-ai-in-healthcare.

57. Hasanzadeh F, Josephson CB, Waters G, Adedinsewo D, Azizi Z, White JA. Bias recognition and mitigation strategies in artificial intelligence healthcare applications. npj Digit Med. 2025; 8:158. doi:10.1038/s41746-025-01503-7.

58. Santa Clara University, Markkula Center for Applied Ethics. Opening the black box: how artificial intelligence increases race-based disparities in medicine. 2025 [cited 2026 Jul 4]. Available from: https://www.scu.edu/ethics/healthcare-ethics-blog/opening-the-black-box-how-artificial-intelligence-increases-race-based-disparities-in-medicine.

59. Stanford Medicine. Clinical AI has boomed [Stanford-Harvard state of clinical AI 2026]. 2026 [cited 2026 Jul 4]. Available from: https://medicine.stanford.edu/news/current-news/standard-news/clinical-ai-has-boomed.html.

60. Systematic review of cost effectiveness and budget impact of artificial intelligence in healthcare. npj Digit Med. 2025. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12381244/ [cited 2026 Jul 4].

61. Trilateral Research. EU AI Act compliance timeline: key dates for 2025-2027 by risk tier. 2025 [cited 2026 Jul 4]. Available from: https://trilateralresearch.com/responsible-ai/eu-ai-act-implementation-timeline-mapping-your-models-to-the-new-risk-tiers.

62. Bignami EG, Russo M, Semeraro F, Bellini V. Balancing innovation and control: the European Union AI Act in an era of global uncertainty. JMIR AI. 2025. doi:10.2196/75527. PMID: 41166703; PMCID: PMC12574960. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12574960.

63. Bipartisan Policy Center. Paying for AI in U.S. health care. Washington (DC): BPC; 2026 [cited 2026 Jul 4]. Available from: https://bipartisanpolicy.org/issue-brief/paying-for-ai-in-u-s-health-care.

64. Cardiology Innovation. CMS includes AI-ECG technology in 2025 reimbursement schedule. 2024 [cited 2026 Jul 4]. Available from: https://cardiologyinnovation.com/artificial-intelligence/ai-based-diagnosis/cms-includes-ai-ecg-technology-2025-reimbursement-schedule.

65. Sidley Austin LLP. Medicare reimbursement pathway for AI-enabled medical devices considered. 2025 May [cited 2026 Jul 4]. Available from: https://www.sidley.com/en/insights/newsupdates/2025/05/medicare-reimbursement-pathway-for-ai-enabled-medical-devices-considered.

66. Office of the National Coordinator for Health IT / ASTP. Data liquidity, affordability, and access: the history and growth of TEFCA. 2026 [cited 2026 Jul 4]. Available from: https://healthit.gov/resources/data-liquidity-affordability-and-access-the-history-growth-of-tefca.

67. Firely. The state of FHIR in 2025: growing adoption and evolving maturity. 2025 Jun 26 [cited 2026 Jul 4]. Available from: https://fire.ly/blog/the-state-of-fhir-in-2025.

68. JD Supra. HHS crackdown on information blocking: new era of enforcement. 2026 [cited 2026 Jul 4]. Available from: https://www.jdsupra.com/legalnews/hhs-crackdown-on-information-blocking-1730008.

69. Edenlab. Interoperability challenges in medium healthcare organizations. 2025 [cited 2026 Jul 4]. Available from: https://edenlab.io/blog/interoperability-strategy-healthcare-organizations.

70. Guillot C. Malpractice in the digital age: EMRs, AI and telehealth fuel liability concerns. Risk & Insurance; 2025 Sep 12 [cited 2026 Jul 4]. Available from: https://riskandinsurance.com/malpractice-in-the-digital-age-emrs-ai-and-telehealth-fuel-liability-concerns.

71. Harvard Medical School. AI implications for health equity: shaping the future of health care quality and safety. Boston (MA): HMS; 2025 [cited 2026 Jul 4]. Available from: https://learn.hms.harvard.edu/insights/all-insights/ai-implications-health-equity-shaping-future-health-care-quality-and-safety.

72. Hussain SA, Bresnahan M, Zhuang J. The bias algorithm: how AI in healthcare exacerbates ethnic and racial disparities - a scoping review. Ethn Health. 2025. doi:10.1080/13557858.2024.2422848. PMID: 39488857. Available from: https://pubmed.ncbi.nlm.nih.gov/39488857.

73. Can AI bridge or widen maternal health inequities? Public Health Chall. 2025. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12445195/ [cited 2026 Jul 4].

74. Perez K, Wisniewski D, Ari A, Lee K, Lieneck C, Ramamonjiarivelo Z. Investigation into application of AI and telemedicine in rural communities: a systematic literature review. Healthcare (Basel). 2025; 13(3):324. doi:10.3390/healthcare13030324. PMID: 39942513; PMCID: PMC11816903. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC11816903.

75. OpenAI. Pioneering an AI clinical copilot with Penda Health. 2025 Jul 22 [cited 2026 Jul 4]. Available from: https://openai.com/index/ai-clinical-copilot-penda-health.

76. European Commission. European Health Data Space (EHDS) Regulation. Brussels: EC; 2025 [cited 2026 Jul 4]. Available from: https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en.

77. Jimerson Birr. Healthcare AI regulation 2026: new compliance requirements every provider must know. 2026 Feb [cited 2026 Jul 4]. Available from: https://www.jimersonfirm.com/blog/2026/02/healthcare-ai-regulation-2025-new-compliance-requirements-every-provider-must-know.

78. ecancer. AI tech should augment physician decision-making, not replace it [ACP policy, Annals of Internal Medicine]. 2024 [cited 2026 Jul 4]. Available from: https://ecancer.org/en/news/24837-ai-tech-should-augment-physician-decision-making-not-replace-it.

79. athenahealth. Embracing the future of AI in healthcare: physicians' insights [Physician Sentiment Survey 2024]. 2024 [cited 2026 Jul 4]. Available from: https://www.athenahealth.com/resources/blog/ai-in-healthcare-physician-sentiments.

80. Association of American Medical Colleges. Artificial intelligence competencies across the learning continuum. Washington (DC): AAMC;

2025 [cited 2026 Jul 4]. Available from: https://www.aamc.org/about-us/mission-areas/medical-education/ai-competencies.

81. Stanford Medicine Magazine. AI makes its way into medical school curriculum. 2025 [cited 2026 Jul 4]. Available from: https://stanmed.stanford.edu/ai-medical-school-curriculum.

82. Nuance Communications. Northwestern Medicine deploys DAX Copilot embedded in Epic within its enterprise to improve patient and physician experiences. PR Newswire; 2024 Aug 15 [cited 2026 Jul 4]. Available from: https://www.prnewswire.com/news-releases/northwestern-medicine-deploys-dax-copilot-embedded-in-epic-within-its-enterprise-to-improve-patient-and-physician-experiences-302222444.html.

83. William Blair. The evolution of value-based care. 2025 [cited 2026 Jul 4]. Available from: https://www.williamblair.com/Insights/The-Evolution-of-Value-Based-Care.

84. Bruce G. 700 lives, $100M saved: healthcare AI ROI in ’25. Becker’s Hospital Review; 2025 Dec 22 [cited 2026 Jul 4]. Available from: https://www.beckershospitalreview.com/healthcare-information-technology/ai/700-lives-100m-saved-healthcare-ai-roi-in-25.

85. Centers for Medicare & Medicaid Services. Lessons from CMS’ Acute Hospital Care at Home initiative. Baltimore (MD): CMS; 2024 [cited 2026 Jul 4]. Available from: https://www.cms.gov/blog/lessons-cms-acute-hospital-care-home-initiative.

86. Hospital at Home Users Group. Summary of U.S. evidence on the Hospital at Home model. 2025 [cited 2026 Jul 4]. Available from: https://www.hahusersgroup.org/wp-content/uploads/2025/09/250821-HaH-Evidence-Two-Pager-V1-1.pdf.

87. Hospital at Home Users Group. Summary of U.S. evidence on the Hospital at Home model [Advocate Health data; same source as reference 86]. 2025 [cited 2026 Jul 4]. Available from: https://www.hahusersgroup.org/wp-

content/uploads/2025/09/250821-HaH-Evidence-Two-Pager-V1-1.pdf.

88. Scheid MR, Friedmann B, Oppenheim M, Hirsch JS, Zanos TP. Development and validation of a clinical wearable deep learning based continuous inhospital deterioration prediction model. Nat Commun. 2025; 16. doi:10.1038/s41467-025-65219-8.

89. Healthcare IT Leaders. AI jobs in healthcare. 2023 [cited 2026 Jul 4]. Available from: https://www.healthcareitleaders.com/blog/ai-jobs-in-healthcare.

90. Microsoft. A year of DAX Copilot: healthcare innovation that refocuses on the clinician-patient connection. Official Microsoft Blog; 2024 Sep 26 [cited 2026 Jul 4]. Available from: https://blogs.microsoft.com/blog/2024/09/26/a-year-of-dax-copilot-healthcare-innovation-that-refocuses-on-the-clinician-patient-connection.

91. Shroff D. AI-designed drugs hit 90% Phase 1 success rate in trials. 2 Minute Medicine; 2026 Feb 19 [cited 2026 Jul 4]. Available from: https://www.2minutemedicine.com/ai-designed-drugs-hit-90-phase-i-success-rate-in-trials.

92. Dermawan D, Alotaiq N. From lab to clinic: how artificial intelligence (AI) is reshaping drug discovery. Pharmaceuticals (Basel). 2025. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12298131/ [cited 2026 Jul 4].

93. Sher R. Nobel Prize celebrates AI's role in protein structure innovation. Drug Target Review; 2024 Dec 9 [cited 2026 Jul 4]. Available from: https://www.drugtargetreview.com/article/154645/nobel-prize-celebrates-ais-role-in-protein-structure-innovation.

94. PharmiWeb. Healthcare digital twin market analysis 2030: technology, applications, impact and forecast. 2026 Feb 10 [cited 2026 Jul 4]. Available from: https://www.pharmiweb.com/press-release/2026-02-10/healthcare-digital-twin-market-analysis-2030-technology-applications-and-impact-and-forecast.

95. Rawat A. Digital twins in healthcare: transforming patient care. Appinventiv; 2023 [cited 2026 Jul 4]. Available from: https://appinventiv.com/blog/digital-twins-in-healthcare.

96. Ekinci O. Building the next layer of clinical AI with hybrid intelligence [Epic Systems and Microsoft Research CoMET]. Roche Diagnostics; 2026 [cited 2026 Jul 4]. Available from: https://diagnostics.roche.com/global/en/healthcare-transformers/article/clinical-ai-hybrid-intelligence.html.

97. Karthikesalingam A, Natarajan V. AMIE: a research AI system for diagnostic medical reasoning and conversations. Google Research; 2024 [cited 2026 Jul 4]. Available from: https://research.google/blog/amie-a-research-ai-system-for-diagnostic-medical-reasoning-and-conversations.

98. Crescendo.ai. The latest AI news and breakthroughs in healthcare and medical [Microsoft MAI-DxO]. 2025 [cited 2026 Jul 4]. Available from: https://www.crescendo.ai/news/ai-in-healthcare-news.

99. Kellton. Gartner predictions on AI agents in healthcare. LinkedIn; 2025 [cited 2026 Jul 4]. Available from: https://www.linkedin.com/posts/kellton_generativeai-agenticai-digitalhealth-activity-7350507795718336512-uJwk.

100. Carmouche D. Radically rethinking primary care: continuous and connected AI-enabled access to maximize health outcomes. Lumeris; 2026 [cited 2026 Jul 4]. Available from: https://www.lumeris.com/in-practice/radically-rethinking-primary-care-continuous-and-connected-ai-enabled-access-to-maximize-health-outcomes.

101. European Commission. Estonia 2025 Digital Decade country report. Brussels: EC; 2025 [cited 2026 Jul 4]. Available from: https://digital-strategy.ec.europa.eu/en/factpages/estonia-2025-digital-decade-country-report.

102. Kriisa K. Digital health: a democratic right to control our own data. e-Estonia; 2025 Sep 17 [cited 2026 Jul 4]. Available from: https://e-estonia.com/digital-health-right-to-control-our-own-data.

AI in Emergency Medicine: What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution, First Edition. Harvey Castro.
© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

9
Leading the AI Shift: What Every CEO, CIO, and CMO Must Know

I've read plenty of chapters that tell CEOs and CIOs what to do about AI. Most of them are written by people who have never run anything.

I ran emergency rooms. I ran a freestanding ER network with more than 400 employees across multiple locations. I've managed budgets, hired and fired staff, negotiated with insurance companies, dealt with malpractice attorneys, and sat across the table from board members who wanted to know why the numbers looked the way they did. I've been the person in the chair when the decision must be made by morning.

So when I tell you that AI governance in healthcare isn't a technology problem, it's a leadership problem, I'm not quoting a white paper. I'm telling you what I learned by living it.

Here's what most consultants won't say: the hardest part of deploying AI in a health system isn't the algorithm. It's the people. It's the chief medical officer who's two years from retirement and doesn't want to be the person who approved the tool that missed something. It's the IT director who's already running a dozen legacy systems and can't absorb another integration project. It's the frontline nurses who've been burned by the last three "workflow improvements" that added 30 minutes to their shift. These aren't irrational objections. They're rational responses to real institutional history.

I learned this at Trusted ER. We built the network from scratch, which meant we could design workflows without inheriting decades of accumulated process debt. But even in a greenfield environment, the friction was real. Every new tool I introduced required convincing clinicians that it would help them, not surveil them. Every dashboard I built for the C-suite (chief executive officer, chief information officer, and chief marketing officer) had to be useful to the nurse on the floor or it was dead

on arrival. Every technology decision had a human cost and a human benefit, and the job of leadership was to be honest about both.

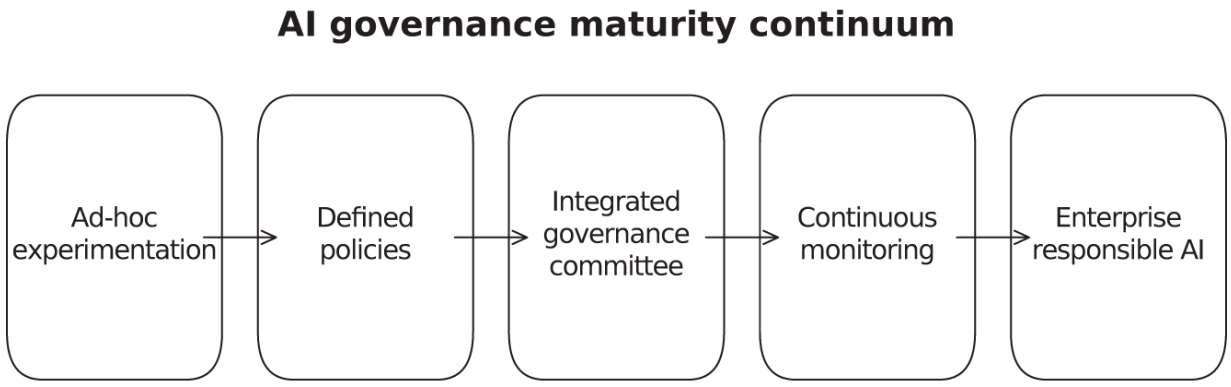
That experience is what I bring to this chapter. The evidence on healthcare AI investment, governance frameworks, and deployment strategy is substantial, and it's covered in detail in the pages that follow. But the evidence doesn't implement itself. Someone must lead the implementation. Someone must build the governance committee, set the budget, manage the vendor relationship, monitor for bias, and answer to the board when something goes wrong.

This chapter is for that someone. Not the theoretical leader. The one who's in the room.

Executive Summary

The single most reliable predictor of whether a healthcare AI program succeeds or fails is not the algorithm, the vendor, or the budget. It's the quality of executive leadership steering the initiative. Healthcare is entering a decisive inflection point. The global AI in healthcare market, valued at approximately $26.57 billion in 2024, is projected to reach $505.59 billion by 2033, a compound annual growth rate of roughly 38.9 percent [1]. That trajectory represents not a gradual evolution but a structural reshaping of medicine itself, one that will reward organizations whose C-suites are prepared and penalize those that are not.

Preparation remains uneven. While 74 percent of U.S. hospitals now report that multiple entities are accountable for evaluating predictive AI and 66 percent have established a specific committee or task force dedicated to AI governance [2], the same figures reveal that most health systems still lack complete strategies for scaling, governing, and measuring AI across the enterprise. A Deloitte survey of 100 healthcare technology executives found that a clear "AI divide" is emerging between early adopters (59 percent of whom expect cost savings exceeding 20 percent within two to three years) and organizations still watching from the sidelines, where only 13 percent project equivalent savings [3].



Conceptual diagram based on Joint Commission/CHAI Responsible Use of AI guidance (Sept 2025).

Figure 9.1 AI adoption S-curve: pilot to early adoption to majority to optimization

The cost of inaction is rising. The American Medical Association documented that physician AI adoption nearly doubled in a single year, jumping from 38 percent to 66 percent between 2023 and 2024 [4]. Healthcare AI investment exceeded $18 billion in 2025, representing 46 percent of all healthcare investment, the highest AI share ever recorded in the sector [5]. The FDA has authorized more than 1,250 AI-enabled medical devices [6], creating an urgent governance imperative that most boards and executive teams are not yet equipped to fulfill (see Figure 9.1).

This chapter gives every CEO, CIO, and CMO the strategic, operational, and clinical frameworks needed to lead, not just manage, the AI shift in their organizations. It draws on peer-reviewed evidence, government data, and documented case studies from named health systems and named executives who have already begun the work. The chapter moves from the strategic case for action through real-world examples, risks and limits, the regulatory field, ethical obligations, a concrete implementation framework, and a forward view of where this transformation leads. The goal is not to sketch the future in broad strokes but to equip senior leaders with the specific knowledge they need to act now.

Current Evidence

Executives face unprecedented pressure to harness AI capabilities while managing organizational risk. This tension defines the current landscape of C-suite decision-making around artificial intelligence deployment.

The State of C-Suite AI Adoption

The empirical picture of where American health systems stand in their AI journey is at once encouraging and sobering. AI use in healthcare is accelerating faster than any other technology deployment in the sector's recent history.

Federal government data from the U.S. Office of the Assistant Secretary for Technology Policy gives the most authoritative baseline. In a September 2025 data brief based on surveys covering more than 2,000 hospitals, ASTP found that 71 percent of nonfederal acute-care hospitals reported using predictive AI integrated with their electronic health record in 2024, a statistically significant increase from 66 percent in 2023 [7]. The figures reveal sharp stratification: 96 percent of large hospitals (more than 400 beds) used predictive AI, compared with 59 percent of small hospitals (fewer than 100 beds). System-affiliated hospitals used predictive AI at 86 percent versus only 37 percent for independent facilities [7]. These gaps signal that the "AI divide" Deloitte identifies at the organizational level is also a structural divide across the care delivery field. The largest systems are pulling ahead rapidly while smaller and independent organizations risk falling behind.

Governance infrastructure is building, though unevenly. In 2024, 74 percent of hospitals had multiple entities accountable for evaluating predictive AI, and 66 percent had established a specific AI committee or task force [8]. Among those evaluating their tools, 82 percent assessed for accuracy, 74 percent assessed for bias, and 79 percent conducted post-implementation monitoring [7]. Roughly one in four hospitals was still not systematically evaluating bias in deployed AI systems, a gap with direct patient safety and legal implications.

On the physician side, the AMA's 2024 Augmented Intelligence Research survey documented the near-doubling of AI adoption among practicing physicians, from 38 percent in 2023 to 66 percent in 2024. The AMA described this jump as "unusually fast" [4]. The top AI opportunity identified by physicians, cited by 57 percent, was reducing administrative burden through automation, signaling that physician-facing AI tools for documentation and workflow are the clearest near-term adoption pathway [4]. Critically, 86 percent of physicians surveyed said they want to be

“responsible or consulted” in AI tool implementation at their organizations [9], a finding with direct implications for CMO-led governance.

CEO Strategic Frameworks

Understanding these adoption trends requires examining how different organizational roles approach AI implementation and the strategic implications for executives.

Three major consulting firms (McKinsey, BCG, and Deloitte) have converged on similar frameworks for C-suite AI strategy, and their guidance is worth examining in sequence because it maps the full arc from vision to execution.

McKinsey’s 2025 analysis argues that healthcare AI is shifting from fragmented point solutions toward federated, modular architectures and clinical data foundries. The firm frames the strategic choice for health system leaders as binary: remain passive buyers and users of third-party tools or become active shapers of an integrated healthcare AI architecture [10]. McKinsey also reports that nearly 80 percent of organizations surveyed draw on AI in at least one function, and 92 percent of executives plan to increase investment over the next three years [11]. This isn’t an incremental optimization story. It’s a redesign of how healthcare organizations operate.

BCG’s December 2025 CEO roadmap for agentic AI—systems capable of independently planning, making decisions, and taking actions toward defined goals with minimal human intervention introduces a “Now, Next, and Always” framework built around five strategic principles:

- Redesigning work from scratch using outcome-driven processes
- Building a shared AI platform and context fabric
- Preparing for a radically different operating model
- Concentrating approximately 80 percent of transformation effort on end-to-end workflow redesign
- Following the “10-20-70 rule,” which ensures the business, not IT, leads transformation [12]

BCG observed that AI-native firms are already achieving 25 to 35 times more revenue per employee than traditional peers, a competitive threat that health system leaders can't dismiss as irrelevant to their market [12].

Deloitte's February 2026 report, based on a September 2025 survey of 100 healthcare technology executives, found that 85 percent of healthcare leaders plan to increase investment in agentic AI over the next two to three years, and 61 percent are already building or implementing initiatives [3]. Deloitte's five-part activation strategy for healthcare organizations stresses a "problem-first, technology-second" approach: identify operational, clinical, or financial problems with enough scale and clarity to warrant AI investment and then build infrastructure, establish an operational foundation, run pilots, and develop production-ready solutions in an iterative cycle [13].

Every one of these frameworks maps to decisions I faced running Trusted ER. The consulting language is different, but the underlying question is the same one I asked every time: does this tool solve a real operational problem, or does it solve the problem of appearing innovative?

The Investment Picture

The financial case for C-suite engagement rests on three converging data streams:

- Escalating external investment that is reshaping the vendor field
- Rising internal health system spending
- Early but compelling ROI evidence

AI investment in healthcare exceeded $18 billion in 2025, representing 46 percent of all healthcare investment, according to Silicon Valley Bank's 17th Healthcare Investments and Exits Report [5]. Deals over $300 million accounted for 40 percent of total healthcare AI spending in 2025, up from 29 percent in 2023, and the total number of healthcare AI mega-deals in 2025 surpassed the sector's overall peak investment year of 2021 [5]. AI-enabled healthcare startups raised rounds 83 percent larger on average than non-AI peers in the same period [14].

At the health system level, internal AI budgets nearly tripled year over year to reach approximately $1.4 billion in 2025. Three categories dominate that spending:

- Ambient clinical documentation ($600 million)
- Coding and billing automation ($450 million)
- Patient engagement and prior authorization (which grew 10-fold year over year) [15]

For 60 percent of healthcare organizations, AI budgets are now growing faster than overall IT budgets, a signal that AI has moved from experimental to strategic [15].

A Menlo Ventures survey of more than 700 health system, outpatient, payer, and life sciences leaders found that 22 percent of healthcare organizations have implemented domain-specific AI tools, a sevenfold increase over 2024 and a 10-fold increase over 2023 [16]. Healthcare AI adoption is accelerating at two times the pace of the broader economy. Medical documentation and back-office revenue cycle management together account for 60 percent of all healthcare IT spending, representing what Menlo Ventures estimates as a $38 billion opportunity for AI to build on established workflows [16].

FDA Authorization and Governance Urgency

The regulatory picture adds a dimension to the governance imperative that many C-suite leaders haven't fully internalized. As of July 2025, the FDA had authorized more than 1,250 AI-enabled medical devices for marketing in the United States, with radiology accounting for approximately 75 to 80 percent of authorized devices and cardiovascular applications accounting for approximately 10 percent [6]. The pace of authorization has accelerated sharply: from roughly 690 devices at the end of 2023 to more than 1,250 within 18 months [6].

Each of those authorized devices represents a potential deployment in a clinical workflow, carrying a governance obligation. The FDA's Predetermined Change Control Plan (PCCP) framework, finalized in December 2024, allows manufacturers to implement pre-approved

algorithmic modifications without new premarket review [17], meaning AI systems can change after deployment. Health systems that don't have active post-deployment monitoring frameworks will be managing systems that may perform differently from how they were originally validated.

Why Clinical Leadership Predicts Success

The single most consequential finding in the recent clinical AI literature comes from a 2025 integrative analysis of 105 clinical studies published in *npj Digital Medicine*. Every CMO must understand it, and every CEO must operationalize it. Analyzing 96 randomized controlled trials and nine non-RCT studies, researchers found that AI deployments led by clinicians were significantly more likely to achieve meaningful clinical impact than those led by technologists. Clinician-led studies showed significant AI impact in 94 percent of cases, compared with 60 percent for technologist-led deployments. Logistic regression found that clinical leadership was associated with 7.79 times higher odds of significant AI impact after controlling for AI type, clinical setting, and team size [18]. Among RCTs only, the odds ratio climbed to 19.898 [18]. The implication is unambiguous: C-suite leaders who structure AI programs without clinical leadership embedded in every layer, from governance to deployment, are systematically cutting their probability of success.

Real-World Case Studies

The organizations profiled in this section represent the leading edge of healthcare AI deployment in the United States. Each has moved past pilot programs into enterprise-scale operations, generating outcome data that smaller implementations can't produce. Their experiences, including both their successes and their unexpected complications, offer the closest thing to a field manual that exists today for health system leaders evaluating AI investment.

Kaiser Permanente: Scaling Ambient AI Across 40 Hospitals

Kaiser Permanente offers the most extensively documented example of large-scale ambient AI deployment in the United States. In late 2023, the

organization created the role of vice president of artificial intelligence and emerging technologies and hired Dr. Daniel Yang, MD, a practicing clinician and former Gordon and Betty Moore Foundation AI governance advisor, to fill it. Dr. Yang reports to Chief Medical Officer Dr. Andrew Bindman and oversees quality oversight for all AI applications across Kaiser's 40 hospitals and 600+ medical offices, spanning clinical operations, research, education, and administrative functions [19].

Under Dr. Yang's leadership, Kaiser Permanente deployed Abridge, an AI-enabled clinical documentation tool, across all 40 hospitals and 600+ medical offices. By September 2024, clinicians had used the tool to record more than 2 million patient interactions, what Kaiser describes as the largest implementation of ambient listening technology in the United States [19]. Physicians reported more meaningful patient interactions, and many began dictating exam findings aloud for the first time, a behavioral change suggesting the technology is genuinely altering the dynamics of the clinical encounter rather than simply automating a documentation task [19].

Kaiser has also codified its governance philosophy in a set of seven principles for responsible AI use, articulated by Dr. Yang:

- Transparency
- Accountability
- Safety
- Equity
- Quality
- Privacy
- Clinician oversight

Kaiser's Advance Alert Monitor, an AI tool estimated to save approximately 500 lives annually, exemplifies these principles in practice [20]. The structural lesson from Kaiser is that scale is achievable when governance and deployment are designed together from the outset, with a clinician holding organizational authority and direct CMO reporting.

Cleveland Clinic: First Chief AI Officer and Strategic Partnerships

Cleveland Clinic represents a different but complementary model, one built around dedicated executive AI leadership and strategic external partnership. In August 2024, the organization appointed Ben Shahshahani, PhD, as its inaugural chief artificial intelligence officer. Shahshahani brought more than 20 years of experience in AI and machine learning from SiriusXM/Pandora, Verizon Media/Yahoo, and Google. He reports to Chief Digital Officer Rohit Chandra, PhD, and his mandate encompasses enterprise AI strategy, patient care improvement, caregiver experience, operational efficiencies, safety, regulatory compliance, and data security [21].

On the clinical documentation side, Cleveland Clinic CIO Sarah Hatchett led a deliberate selection process, conducting five pilots over six months before choosing Ambience Healthcare's AI Scribe. The outcome was compelling: within 15 weeks of rollout, more than 4,000 of 6,000 eligible physicians and advanced practice providers were actively using the system. More than 1 million patient encounters were documented. Note-writing and review time fell by 2 minutes per appointment and 14 minutes per day [22]. CDO Rohit Chandra stated: "The promise of AI in healthcare is that it will enable us to care for patients with a higher level of safety and quality, and a better patient and caregiver experience" [22].

Cleveland Clinic also built a Virtual Command Center in partnership with Palantir that draws on AI to create a real-time overview of hospital status, enabling 24-hour patient admission forecasting and producing a 7 percent increase in hospital transfer rates to the main campus [23]. In March 2025, Cleveland Clinic and G42, a global AI firm headquartered in Abu Dhabi, announced a strategic collaboration creating a joint task force to evaluate and accelerate AI projects for patient care, medical research, and operational efficiency globally [24]. Cleveland Clinic CEO Tom Mihaljevic, MD, stated: "Artificial intelligence offers a tremendous opportunity to continue to advance and fulfill our mission of caring for life, researching for health and educating those who serve" [24]. (In this context, the increase reflects the system's improved ability to identify patients at satellite facilities who genuinely needed main-campus resources,

transferring them earlier rather than managing them in settings without adequate specialty support. The metric indicates better triage, not increased hospitalization.)

NYC Health + Hospitals and WashU Medicine/BJC: Public Systems and Academic Centers

NYC Health + Hospitals, the nation's largest municipal health system with more than 43,000 employees and 70+ patient care locations, appointed Divya Pathak, MS, MBA, as its inaugural vice president and chief data and artificial intelligence officer, effective March 2025. Pathak previously served as the system's chief data officer and brought prior experience from Mayo Clinic and IBM Research. CEO Mitchell Katz, MD, personally cited her in the announcement, a clear signal of direct CEO sponsorship for the AI and data transformation agenda [25]. The appointment shows what's possible for publicly funded safety-net systems: AI leadership infrastructure is not exclusively the province of large academic medical centers or well-capitalized private systems.

In St. Louis, Washington University School of Medicine and BJC Health System launched a joint Center for Health AI in 2024, led by inaugural Chief Health AI Officer Philip R.O. Payne, PhD. The operating team includes Paul Scheel, MD, CEO of Washington University Physicians, and Chris Miller, MD, chief clinical officer of BJC, a governance structure that explicitly bridges the academic and health system enterprises. Clinical AI initiatives include tools for high-risk patient flagging for sepsis and chemotherapy complications, AI-assisted medical imaging for tumor detection, surgical blood loss prediction, and ambient documentation to reduce clinician burnout [26]. The center also provides immersive AI training for medical students and residents, treating workforce development as a core component of the AI mission from the outset [26].

These four case studies share structural features that C-suite leaders can learn from:

- Named executive accountability at the vice president or C-suite level
- Clinician involvement in governance from day one
- A deliberate vendor evaluation process with pilot comparisons

- Integration of AI strategy with the broader organizational mission
- Willingness to partner externally when internal capability doesn't yet exist

Risks and Limitations

The scale of healthcare AI deployment described in the previous section brings proportional risk. When AI systems fail at enterprise scale, the consequences are measured in patient harm, regulatory action, and billions of dollars in losses. The cases that follow are not edge cases. They are the most prominent, well-documented failures in healthcare AI to date, and each carries lessons that any organization deploying these tools must internalize.

High-Profile AI Failures and What They Teach

The history of healthcare AI contains cautionary cases that every C-suite leader must understand. Not to discourage investment, but to guard against the specific failure modes that have cost billions of dollars and, in some cases, harmed patients.

IBM Watson for Oncology IBM Watson for Oncology is the most-studied failure. After consuming more than $5 billion in acquisitions and development, IBM sold its health division in 2022 for approximately $1 billion, a documented loss of roughly $4 billion [27]. Real-world concordance with expert oncologists ranged from 12 percent for gastric cancer in Chinese hospitals to 96 percent at institutions using similar clinical guidelines. MD Anderson Cancer Center spent approximately $62 million before halting its Watson oncology project; internal IBM documents reviewed by STAT described "unsafe and incorrect" treatment recommendations. Denmark's national cancer center found only 33 percent concordance and rejected the system [27]. The root cause: Watson was trained primarily on hypothetical synthetic cases rather than real patient data from diverse clinical settings.

Optum The Optum racial bias case illustrates a different failure mode, one with direct legal and ethical consequences. A study published in *Science* in

2019 found that a widely deployed commercial risk prediction algorithm used by health systems across the United States systematically discriminated against Black patients. The algorithm used healthcare costs as a proxy for health needs. Because Black patients spent approximately $1,800 less annually than equally ill white patients due to systemic access barriers, the algorithm directed extra care toward healthier white patients. Black patients categorized by the algorithm as equally needing care as white patients had 48,772 more additional chronic conditions across the population studied. Correcting the bias increased Black patient inclusion in care management programs from 17.7 percent to 46.5 percent [28]. Optum then committed to revising the algorithm. The lesson is plain: AI systems embedded in care management workflows can encode and amplify pre-existing health inequities unless bias evaluation is treated as a governance requirement, not an optional enhancement.

Epic's sepsis prediction model, deployed across hundreds of U.S. hospitals, showed materially different real-world performance from vendor claims. External validation at Michigan Medicine found an area under the curve of 0.63 versus Epic's claimed 0.76 to 0.83, with sensitivity of only 33 percent at recommended thresholds. The model generated alerts for 18 percent of all hospitalized patients while missing 67 percent of actual sepsis cases, requiring clinicians to evaluate 109 flagged patients to find one who actually required sepsis intervention [29]. Alert fatigue was the result, a well-documented failure mode in which clinicians learn to ignore warnings because the signal-to-noise ratio is too poor (positive predictive value under 1 percent). I've seen this firsthand in the ER: when every alert fires, none of them matter anymore.

Cybersecurity: The Change Healthcare Precedent

On February 21, 2024, the ransomware group BlackCat attacked Change Healthcare, a UnitedHealth Group subsidiary that processes approximately half of all medical claims in the United States, serving roughly 900,000 physicians, 33,000 pharmacies, 5,500 hospitals, and 600 laboratories. The attackers gained access through compromised credentials for a Citrix remote access server that did not have multifactor authentication enabled. UnitedHealth Group CEO Andrew Witty confirmed in congressional testimony that $22 million in bitcoin was paid as ransom [30]. The breach

ultimately exposed the protected health information of an estimated 190 million individuals, the largest healthcare data breach ever reported, surpassing the prior record of 78.8 million from the Anthem breach in 2015 [30]. The total anticipated cost was revised to $2.87 billion in UnitedHealth Group's Q3 2024 earnings report [30].

An AHA survey of nearly 1,000 hospitals conducted in March 2024 found that 74 percent reported direct patient care impact from the attack, 94 percent reported financial impact, and 33 percent reported that the attack disrupted more than half of their revenue [31]. The AHA characterized it as "the most significant and consequential incident of its kind against the U.S. healthcare system" [31]. The structural lesson for C-suite leaders deploying AI is this: AI systems are not isolated technology tools. They're nodes in an interconnected digital infrastructure, and concentration of critical functions in single-vendor dependencies creates systemic national-scale vulnerability.

Change Management Gaps

The Kotter model survey data referenced in European hospital digital transformation research reveals that more than half of hospitals use no formal change management framework at all when implementing technology [32]. Among those that do, Kotter's eight-step model is most frequently cited, but its top-down orientation can leave frontline staff feeling that change is being done to them rather than with them [32]. This finding has direct implications for AI deployment: without structured change management, organizations invest in tools that clinicians don't adopt, governance structures that committees don't follow, and strategies that never reach the bedside.

Talent Shortages

The talent challenge facing health system CIOs is compounding. ISC2 projected a global workforce gap of 4.8 million in the cybersecurity industry, with 47 percent of jobs going unfilled [33]. The hiring gap for AI roles specifically is estimated at 50 percent [33]. An Amazon/Gallup survey found that while 27 percent of healthcare workers expressed desire to build technical or digital skills, only 6 percent participated in actual upskilling in those areas [33]. A GlobalData survey found that 43 percent of healthcare professionals cited lack of IT expertise as having delayed digital

transformation in their organizations [33]. These shortages mean that organizations can't simply hire their way to AI capability. They must invest in systematic workforce development alongside deployment.

Build vs. Buy Complexity

A KPMG survey found that only 12 percent of organizations are developing generative AI solutions exclusively in-house [34]. Creating large language models can cost organizations anywhere from $1 million to $100 million, which explains why most organizations adopt a blend of building, acquiring, and partnering [34]. For most health systems, buying from established vendors with demonstrated clinical validation aligns better with organizational capability and risk tolerance. Buying still requires rigorous due diligence on vendor performance claims, as the Watson and Epic sepsis cases demonstrate.

Regulatory and Legal Landscape

The regulatory infrastructure for healthcare AI is being built while the technology is already in clinical use. The FDA, state legislatures, and international bodies are all moving, but at different speeds and with different frameworks. What follows is the current state of that regulatory landscape as of early 2026, with attention to the specific mechanisms that will most directly affect health system AI procurement and deployment decisions.

FDA AI/ML Framework and the PCCP

The FDA's regulatory approach to AI-enabled medical devices has evolved substantially over the past several years and culminated in two major developments in late 2024 and early 2025 that every CIO and CMO must understand.

The FDA's Total Product Lifecycle (TPLC) approach requires oversight from development through post-market monitoring, grounded in 10 Good Machine Learning Practice principles published in 2021 jointly with Health Canada and the UK's MHRA [35]. The framework addresses data

management, transparent model development, and continuous monitoring as foundational requirements [35].

The most consequential recent development is the finalization of the Predetermined Change Control Plan guidance in December 2024. PCCPs allow manufacturers to specify in their initial marketing submission exactly which types of algorithmic modifications they anticipate making and to implement those modifications without submitting new premarket review applications, provided the changes remain within the PCCP's scope [17]. A PCCP must include a description of planned modifications, a modification protocol specifying methodology for developing and validating changes, and an impact assessment. The practical implication for health systems is significant: AI systems they deploy today may be modified by vendors without triggering a new FDA review. Post-deployment monitoring isn't optional. It's clinically and legally necessary.

The FDA also published draft guidance on AI-enabled device software lifecycle management in January 2025 and issued a request for public comment in September 2025 on measuring real-world AI-enabled device performance. A recent study found that approximately 6 percent of AI-enabled devices have faced recalls, often within the first year [36].

ONC HTI-1: 31 Transparency Attributes

The ONC Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing Final Rule (HTI-1), published in January 2024 and effective for implementation from January 2025, created the first federal transparency requirements for AI and predictive algorithms in certified health IT [37]. The rule distinguishes between evidence-based decision support interventions, which require 13 source attributes, and predictive decision support interventions, defined as technology that derives relationships from training data to produce predictions, classifications, recommendations, or evaluations, which require 31 source attributes [37]. These 31 attributes include training datasets, external validation, performance metrics, and fairness and bias evaluation processes. Health IT developers must also apply Intervention Risk Management practices (risk analysis, mitigation, and governance) and submit public summary information via a publicly accessible hyperlink

[37]. The HTI-1 rule establishes a minimum transparency standard for any AI tool integrated with a certified EHR.

CMS Reimbursement: 26 CPT Codes and the WISeR Pilot

As of January 2026, there are 26 CPT codes for clinical AI solutions. Yet only three AI solutions have received Category I CPT codes, which indicate broad clinical adoption and established evidence:

- FFR-CT for cardiovascular risk prediction via CT imaging
- A service analyzing CT scan data to predict cardiovascular risk
- A tool using augmented imaging to detect diabetic retinopathy [38]

All other AI solutions carry Category III codes (temporary codes used for data collection purposes), which means reimbursement varies by Medicare Administrative Contractor jurisdiction and is not uniformly available [38].

CMS launched the Wasteful and Inappropriate Service Reduction (WISeR) model on January 1, 2026, in six states: Arizona, Ohio, Oklahoma, New Jersey, Texas, and Washington. Running through December 31, 2031, this pilot allows AI to support prior-authorization decisions for certain medical services in Original Medicare. Initial services subject to AI-assisted review include electrical nerve stimulator implants, epidural steroid injections, and percutaneous vertebral augmentation [39]. In December 2025, CMS also announced the ACCESS model, a voluntary initiative providing recurring monthly payments to organizations managing patients with chronic conditions, an explicit alignment of value-based care payment with AI-enabled care delivery [40].

State Laws: California and Colorado

California enacted Senate Bill 1120 on September 28, 2024, effective January 1, 2025, governing AI used in use review and management by healthcare service plans. Its core provision is unambiguous: no AI, algorithm, or software tool may deny, delay, or modify healthcare services based on medical necessity determinations. Only a licensed physician or healthcare professional may make that determination. Utilization review

decisions must be based on an individual patient's own records, not "solely on a group dataset" [41].

Colorado's Senate Bill 24-205, enacted in May 2024 and effective June 30, 2026, is described as the first U.S. law to impose full, enforceable AI governance obligations on private-sector organizations. The law applies to developers and deployers of "high-risk AI systems" making or substantially influencing consequential decisions in healthcare, employment, housing, financial services, and insurance. Key obligations include annual impact assessments, risk management programs aligned with NIST or ISO/IEC 42001 standards, consumer notification when AI is involved in consequential decisions, the right of consumers to appeal decisions and correct data, and disclosure to the Colorado Attorney General within 90 days of discovering algorithmic discrimination [42]. More than a dozen additional states are advancing similar legislation, and health system compliance officers must track this rapidly evolving patchwork.

EU AI Act: Extraterritorial Reach

The EU Artificial Intelligence Act entered into force on August 1, 2024, the world's first full legal framework for AI. Its reach extends beyond European borders: any provider or deployer whose AI system outputs occur in the European Union is subject to the regulation. Medical devices, in vitro diagnostics, and software used in healthcare triage are explicitly designated high-risk [43]. Full obligations for stand-alone systems take effect August 2, 2026; for AI embedded in regulated medical devices, August 2, 2027 [43]. Penalties reach €35 million or 7 percent of global annual turnover [43, 73]. U.S. health systems with any AI-enabled products used in EU member states must assess compliance obligations now.

Reading through this regulatory stack (FDA, ONC, CMS, California, Colorado, EU), I'm struck by how much compliance burden now lands on the CMO and general counsel, not just the IT department. When I was building governance frameworks at Trusted ER, we had no playbook for any of this; the organizations that build that legal literacy now will have a real structural advantage.

Fiduciary Duty: The Marchand Precedent

The legal dimension of board AI oversight gained sharper definition with the Delaware Supreme Court's *Marchand v. Barnhill* ruling, which established that boards must actively implement and deploy monitoring systems for "mission critical" risks central to a company's business model, and Caremark fiduciary duties apply most rigorously where oversight failures could threaten enterprise viability [44]. For healthcare boards, AI systems influencing safety-critical diagnostics, treatment recommendations, and staffing decisions are by definition mission critical. The Harvard Medical School Ethics Center notes that 13 percent of S&P 500 companies now have technology committees at the board level, a governance model that healthcare boards should treat as a minimum standard [44].

Ethical Implications

The ethical questions raised by healthcare AI are not abstract. They have concrete policy answers in some jurisdictions and remain unresolved in others. The frameworks described here, from the WHO, the European Union, and national-level bodies, represent the international community's attempt to codify principles that can guide deployment without stifling the innovation that patients need.

International Ethical Frameworks

The ethical architecture for healthcare AI governance has been built primarily through three international frameworks that health system leaders must understand, even if their operational application requires local adaptation.

The World Health Organization published its "Ethics and Governance of Artificial Intelligence for Health" in June 2021, the first global guidance document on AI governance for health, developed through a two-year consultation by an international multidisciplinary expert group representing all WHO regions [45]. WHO Director-General Dr. Tedros Adhanom Ghebreyesus stated: "Like all new technology, artificial intelligence holds enormous potential for improving the health of millions of people around

the world, but like all technology it can also be misused and cause harm" [45]. WHO's six ethical principles are:

- Protecting human autonomy
- Promoting human well-being, safety, and the public interest
- Ensuring transparency, explainability, and intelligibility
- Building responsibility and accountability
- Ensuring inclusiveness and equity
- Promoting AI that is responsive and sustainable [45]

The Organisation for Economic Cooperation and Development (OECD) AI Principles, adopted in May 2019 and endorsed by 46 countries including the United States, Canada, EU member states, and Japan (then endorsed by the G20), have strongly influenced the EU AI Act and national AI strategies [46]. The OECD's 2024 health AI policy brief identifies five categories of required action:

- Enabling oversight within and across countries
- Establishing measurement and public reporting on AI adoption and harms
- Defining requirements and liability for AI in health
- Requiring inclusion of health professionals throughout the AI lifecycle
- Establishing programs to encourage responsible use of authorized AI solutions [46]

The White House Blueprint for an AI Bill of Rights, released in October 2022, identifies five principles applicable to healthcare AI:

- Safe and effective systems
- Algorithmic discrimination protections
- Data privacy
- Notice and explanation

- Human alternatives, consideration, and fallback [47]

The document's "Human Alternatives" principle, which requires that people have access to a human alternative when automated systems are used, directly intersects with the CMS Medicare Advantage policy under 42 CFR § 422.101(c), which specifies that AI cannot "act alone" in coverage determinations affecting medical necessity [47].

AMA Governance and Physician Voice

The AMA released new "Principles for AI Development, Deployment, and Use" in December 2023, covering oversight, transparency, disclosure, bias mitigation, and liability [48]. AMA President Jesse Ehrenfeld, MD, MPH, stated: "The AMA recognizes the immense potential of health care AI. Yet this power comes with ethical considerations and potential risks that demand a proactive and principled approach" [48]. The AMA was explicit that voluntary compliance is insufficient and that mandatory governance frameworks are required [48]. The AMA's 2025 Steps Forward Governance for Augmented Intelligence toolkit, developed with Manatt Health, gives health systems an eight-pillar governance framework covering executive accountability, cross-functional working groups, policy development, intake and vendor evaluation processes, oversight and monitoring, and AI organizational readiness [49].

Bias as Ethical Failure: Lessons from Optum

The Optum racial bias case, published in *Science* in 2019 and documented in detail in the research underpinning this chapter, is the canonical example of how an ethically designed intent (identify high-risk patients for additional care management) can produce discriminatory outcomes at scale when the algorithm's training logic encodes systemic inequity [28]. The algorithm's designers didn't intend to discriminate. The discrimination emerged from the choice of proxy variable: cost as a stand-in for need, in a healthcare system where access itself is unequally distributed. The ethical lesson is plain: good intentions don't substitute for bias auditing. Every health system deploying a commercial AI tool bears ethical and increasingly legal responsibility to validate that tool's performance across demographic subgroups before and after deployment.

ECRI, the global healthcare safety nonprofit, ranked “Risks with AI-enabled health technologies” as the number-one health technology hazard for 2025 [50]. ECRI CEO Marcus Schabacker, MD, PhD, stated: “The promise of artificial intelligence’s capabilities must not distract us from its risks or its ability to harm patients and providers” [50]. For 2026, ECRI elevated “Misuse of AI chatbots in health care” to the number-one hazard, warning that as AI chatbots become normalized in clinical environments, people begin to rely on them without question, precisely when errors are most likely to go undetected [51].

Leadership Implementation Framework

Governance structures for healthcare AI can’t be designed in the abstract. They require clear accountability distributed across specific executive roles, with defined responsibilities that map to the technical, clinical, and financial dimensions of AI deployment. The following framework assigns those responsibilities to the three C-suite roles most directly accountable for AI outcomes in a health system.

Three-Role Framework: CEO, CIO, CMO

Effective AI leadership in a health system requires three distinct executive roles performing distinct functions, not a single “AI champion” attempting to span strategy, infrastructure, and clinical governance simultaneously.

CEO The CEO’s role is strategic and fiduciary. The CEO must define the vision and translate it into organizational commitment: what does it mean for this health system to be an AI-enabled organization in three years, and what will that require in terms of investment, culture, and talent? The CEO is responsible for ensuring that AI strategy is integrated into the broader strategic plan, that resource allocation through the budget process supports the vision, and that the board is properly informed and engaged in AI risk oversight. As the AHA Center for Health Innovation has documented, health systems are elevating technology leadership from IT support functions to enterprise-level value creation roles. That elevation starts with the CEO’s willingness to sponsor it personally [52].

CIO The CIO's role is infrastructure, security, and data governance. The CIO must build and maintain the technical foundation that makes clinical AI possible:

- Cloud infrastructure with appropriate HIPAA-eligible services
- FHIR-based interoperability enabling AI systems to access and write structured clinical data
- Data governance frameworks that ensure training data quality and appropriate access controls
- Cybersecurity architecture incorporating the lessons of Change Healthcare, with particular attention to multifactor authentication for all remote access systems and vendor concentration risk management
- Post-deployment monitoring systems capable of detecting model drift, bias emergence, and anomalous outputs

The NIST AI Risk Management Framework's four core functions (Govern, Map, Measure, and Manage) give a reliable operational structure for the CIO's AI governance mandate [53]. CHIME's 10 AI principles for healthcare, released in July 2024, offer the CIO community a peer-developed governance framework covering patient safety, cybersecurity, regulatory oversight, privacy, equity, and workforce education [54].

CMO The CMO's role is clinical validation, safety, and physician leadership. Drawing directly on the *npj Digital Medicine* finding that clinician-led deployments have 7.79 times higher odds of significant AI impact [18], the CMO must ensure that clinical expertise is embedded in every AI governance and deployment process. This means co-chairing the AI governance committee with a technical leader; requiring pre-deployment clinical validation on local patient data before any AI tool enters a clinical workflow; establishing ongoing post-deployment monitoring for model drift and safety signals; and building the physician champion network that *npj Digital Medicine* research identifies as a key predictor of technology adoption within practice groups [55]. I've found that the physician champion matters more than the technology itself. One trusted colleague who says "this actually works" is worth more than any vendor demo.

AI Governance Structure: Two-Tier Model

Published healthcare AI governance models and the peer-reviewed PPTO framework study from *npj Digital Medicine* both converge on a two-tier committee structure as the operational architecture for health system AI governance.

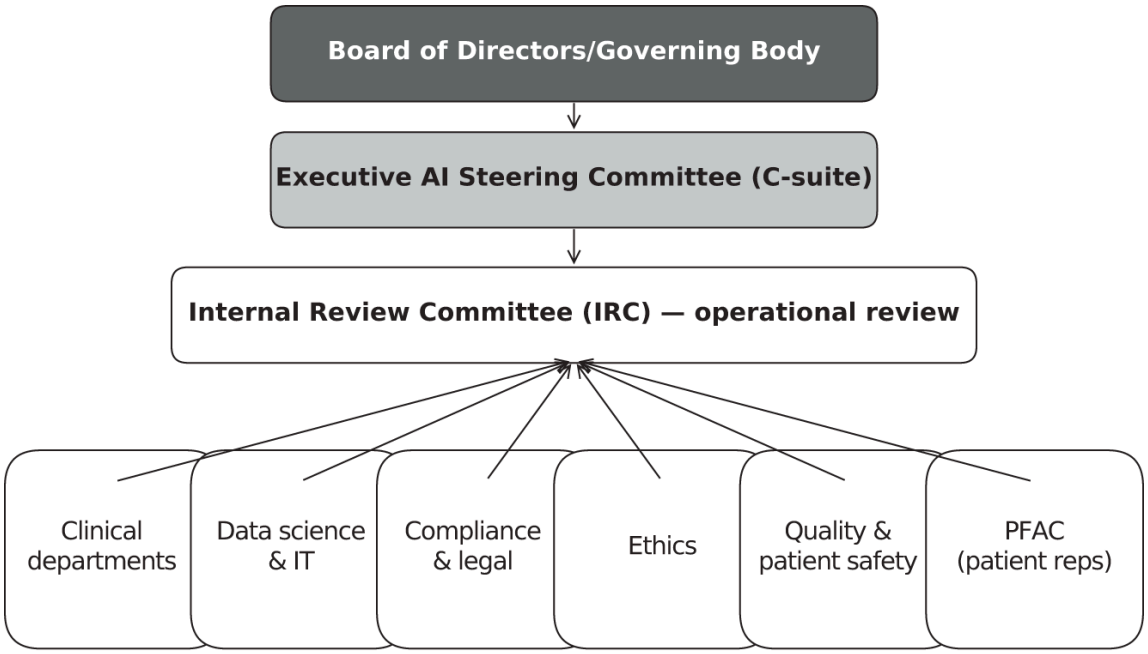
The first tier is the AI Steering Committee, an executive body composed of C-suite leaders (CEO or designee, CIO, CMO, chief information security officer [CISO], chief legal officer [CLO], CFO), responsible for approving the AI strategy, allocating budget, setting the risk tolerance framework, and reporting to the board. Some organizations add external advisors, including patient representatives and ethics experts [56]. The steering committee sets direction and mandates. It doesn't conduct arm's-length technical review.

The second tier is the Internal Review Committee (IRC), a cross-functional body conducting operational review of AI systems at feasibility, technical scope, and Responsible AI stages. IRC membership draws from clinical departments, data science and IT, compliance and legal, ethics, quality and patient safety, and patient representatives engaged through the Patient and Family Advisory Council [57]. The IRC's specific governance responsibilities include:

- Approving AI use cases before deployment
- Reviewing vendor validation evidence against the organization's own patient population
- Conducting bias assessments
- Approving contract language on performance guarantees and fairness requirements
- Establishing post-deployment monitoring protocols with clear accountability for anomalies (see Figure 9.2)

The People, Process, Technology, and Operations (PPTO) framework, studied through action case research at a Canadian hospital system and published in *npj Digital Medicine* in 2025, gives a validated methodology for establishing this governance structure [57]. Key PPTO findings include that:

- The governance committee must be established as a subcommittee of existing leadership (not a new parallel structure).
- Membership must be reviewed annually.



Source: Manuscript Chapter 9 (Internal Review Committee composition).

Figure 9.2 AI governance committee organizational structure

- Individual departments retain responsibility for operational AI activities while the committee maintains systemic oversight.
- Patient advisors must be included from the start [57].

Vendor Evaluation Scorecard

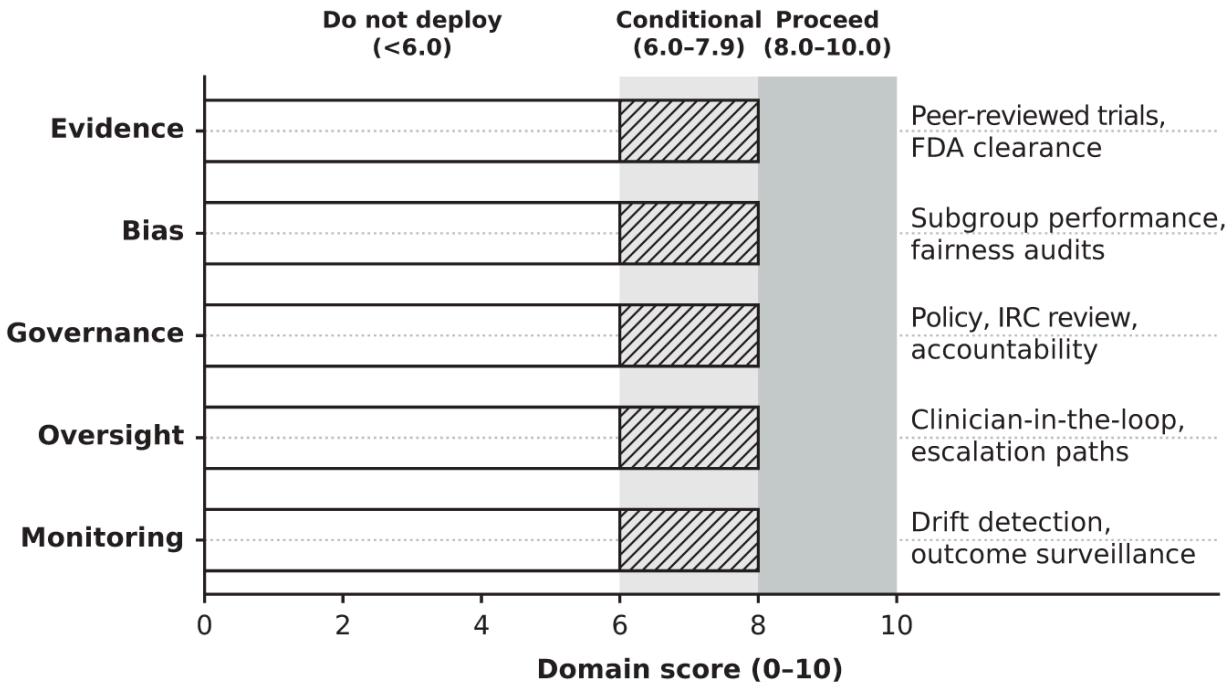
When selecting commercial AI tools, health system leaders need an objective, multidomain evaluation framework that prevents vendor marketing claims from substituting for clinical evidence. The Public Health AI Handbook's Vendor AI Evaluation Toolkit gives a validated six-domain weighted scorecard:

- Technical Validation (25 percent)
- Clinical Safety (25 percent)
- Fairness and Equity (20 percent)

- Privacy and Security (15 percent)
- Workflow Integration (10 percent)
- Business Viability (5 percent) [58]

Each domain is scored 0 to 10. Decision thresholds are:

- 8.0 to 10.0, proceed with deployment
- 6.0 to 7.9, conditional deployment with mitigation plan
- Below 6.0; do not deploy [58] (see Figure 9.3)



Source: Manuscript Chapter 9 — vendor evaluation rubric. Scores entered by Internal Review Committee per domain.

Figure 9.3 AI vendor evaluation scorecard: evidence, bias, governance, oversight, monitoring

For Technical Validation, the minimum acceptable standard requires external, independent validation at three or more sites with full peer-reviewed publication, complete performance reporting across demographic subgroups, and evidence of generalizability to the deploying institution's patient population. A score below 5 out of 10 on this dimension alone is defined as insufficient evidence to deploy [58]. For Privacy and Security,

vendors that refuse to sign a Business Associate Agreement or have no security certifications should receive a score of zero [58].

For contracts, key provisions must include:

- Minimum AUC thresholds with termination rights if unmet
- Quarterly bias audit reports stratified by demographic subgroups with remediation obligations if disparate impact exceeds 10 percent
- Data ownership guarantees
- SLAs covering both technical uptime and clinical performance metrics [58]

Phased Implementation Roadmap

A sustainable healthcare AI implementation follows a phased approach that sequences investment, learning, and scale in a manner aligned with organizational readiness and evidence development.

Phase 1 (Months 1 to 6): Governance Foundation

- Establish the AI Steering Committee and Internal Review Committee with clear charters.
- Conduct a digital maturity and AI readiness assessment using the DiMe Health AI Maturity Model, which evaluates organizations across seven domains from leadership and governance through data infrastructure, workflow integration, and financial sustainability [59].
- Identify three to five high-priority use cases based on clinical need, data availability, and evidence strength; ambient documentation, sepsis detection validation, or revenue cycle automation are high-readiness starting points.

Phase 2 (Months 6 to 18): Controlled Pilots

- Conduct vendor evaluation using the six-domain scorecard.
- Run pilots with defined prespecified metrics, control groups where feasible, and clinical champion co-leadership.

- Validate AI tool performance against local patient population data before broader deployment.
- Use the Consort-AI reporting standards for any pilots intended to generate publishable evidence [60].
- Document outcomes rigorously: timing, workflow impact, patient safety signals, and bias distributions by demographic subgroup.

Phase 3 (Months 18 to 36): Governed Scale

- Deploy validated tools organization-wide with active monitoring, post-deployment drift detection, and a documented incident reporting pathway.
- Integrate AI governance into existing quality and safety infrastructure, including Joint Commission accreditation preparation.
- Build the Analytics Center of Excellence with the eight core roles identified in published frameworks:
  - Clinical AI specialists
  - Data scientists and ML engineers
  - Healthcare data engineers
  - Clinical subject-matter experts
  - Ethics and compliance officers
  - Business analysts
  - Change management specialists
  - UI/UX designers

ROI Measurement

C-suite leaders must build ROI measurement into AI programs from the outset, both to manage board accountability and to prioritize the most valuable use cases.

The strongest peer-reviewed ROI evidence comes from radiology AI. A study published in the *Journal of the American College of Radiology* found

that implementing a multi-application AI platform across 14 solutions in a U.S. stroke management-accredited hospital produced a 451 percent ROI over a five-year period when factoring in labor time reductions, and 791 percent ROI when radiologist time savings were also monetized [61]. The platform generated $3.56 million in revenue for total costs of $1.78 million over five years, corresponding to a return of $4.51 for every dollar invested, or $7.91 when physician time is valued [61].

On the clinical documentation side, the Permanente Medical Group deployed ambient AI scribes across 7,260 physicians and 2.57 million patient encounters over 63 weeks, saving an estimated 15,791 hours of documentation time [62]. In my experience, the burnout reduction alone justifies the investment. When physicians spend fewer minutes on documentation, they spend more minutes with patients. The Mass General Brigham and Emory Healthcare study published in *JAMA Network Open* found a 21.2 percent absolute reduction in burnout prevalence at 84 days among ambient AI scribe users [63, 74]. These burnout reductions translate into direct financial value: physician burnout costs health systems between $800,000 and $1.3 million per physician in recruitment and lost productivity [64].

In revenue cycle management, McKinsey analysis found that AI enablement could cut cost-to-collect by 30 to 60 percent and estimated that a 1 to 2 percentage point reduction in cost-to-collect represents $60 million to $120 million in savings for a health system with $6 billion in patient revenue [65]. Health systems collectively spend more than $140 billion annually on revenue cycle operations, approximately 3 to 4 percent of revenue, making this one of the highest-value AI application categories in the sector [65]. An HFMA-commissioned survey of 556 CFOs and revenue cycle leaders found that health systems using automation already had a cost-to-collect of 3.51 percent versus 3.74 percent for those without automation, a savings of approximately $11.5 million for a $5 billion-revenue system [66].

A rigorous ROI methodology requires:

- SMART objectives and measurable KPIs agreed by stakeholders before deployment

- A 6- to 12-month baseline for data collection before implementation
- Full accounting of direct and indirect costs including infrastructure, data pipeline, model maintenance, and talent
- Projection of benefits across four categories:
  - Direct savings
  - Labor and productivity gains
  - Error and waste reduction
  - Downstream clinical pathway revenue
- Dynamic ongoing ROI measurement, not a one-time calculation at project close [67]

Future Outlook

The investment, governance, and deployment decisions described in this chapter are playing out against a market that is accelerating rapidly. The projections in this section quantify where that market is heading, who is positioned to benefit, and what happens to organizations that delay.

Market Trajectory and the AI Divide

The market projections reinforce the urgency of the governance investments described in this chapter. Grand View Research projects the global AI in healthcare market to grow from $26.57 billion in 2024 to $505.59 billion by 2033 at a CAGR of approximately 38.9 percent [1]. The U.S. market alone is projected to reach $222.9 billion by 2033 [68]. The AI in Hospital Operations segment is projected to grow from $7.51 billion in 2025 to $25.70 billion by 2030 [69]. These projections reflect actual capital flows: healthcare AI investment exceeded $18 billion in 2025, and $10.7 billion in seed-through-growth-stage funding was directed to AI-powered health tech companies through mid-2025, already 24.4 percent higher than all of 2024 [70].

The Deloitte AI divide, the gap between early adopters and watchers, will widen as AI capabilities compound. Organizations that have already validated ambient AI documentation, deployed revenue cycle automation,

and established governance infrastructure will find it progressively easier to add clinical decision support, predictive analytics, and agentic AI. Organizations still in the governance-design or pilot-debate phase in 2026 will be competing against health systems that have been in production operation for two or more years.

The Chief AI Officer as Standard Infrastructure

The appointments documented in this chapter (Ben Shahshahani at Cleveland Clinic, Philip Payne at WashU/BJC, Divya Pathak at NYC Health + Hospitals, and Tommy Ibrahim at Sanford Health) represent an organizational evolution moving from early adoption to standard infrastructure [71]. The AHA Center for Health Innovation noted in February 2026 that traditional IT governance models are being "stretched by the pace of AI and automation moving beyond pilot projects into core clinical and operational workflows, demanding clearer ownership and tighter alignment with organizational strategy" [52]. The chief AI officer role, reporting to either the CEO, CMO, or CDO depending on organizational structure, gives that single point of accountability. Within three to five years, health systems without a dedicated executive AI leadership role will be outliers.

CMS Value-Based Care and AI Alignment

The CMS WISeR pilot and ACCESS model are early signals of a broader reimbursement shift. As value-based care contracts account for a growing share of hospital revenue, AI tools that demonstrably improve outcomes for chronic disease populations, reduce preventable admissions, and support efficient prior authorization will carry reimbursement advantages. The 93 percent of health plan executives who expect AI to add value by automating prior authorizations represent a near-consensus view of where the near-term business case lies [72]. Organizations that build AI capabilities aligned with value-based care metrics will be structurally advantaged as payer mix evolves.

Agentic AI: The Next Transformation Wave

BCG's "Now, Next, and Always" framework points toward a third phase of healthcare AI that is only beginning to emerge. Agentic AI refers to systems

capable of taking multistep actions, reasoning across complex inputs, and coordinating tasks with minimal human intervention [12]. Deloitte's February 2026 report (September 2025 survey) found that 82 percent of healthcare AI early adopters are already prioritizing multi-agent solutions, and 61 percent are building or implementing agentic AI initiatives [3]. At the revenue cycle level, McKinsey describes this as the race to a "touchless revenue cycle," where AI agents autonomously handle claim submission, denial management, authorization tracking, and coding with minimal human review for routine cases [65].

The revenue cycle projections illustrate the magnitude: $60 million to $120 million in savings per year for a health system with $6 billion in patient revenue [65]. Against that backdrop, the governance investments described in this chapter (a steering committee, an internal review committee, a vendor scorecard, a monitoring infrastructure, and a trained AI workforce) aren't overhead costs. They're the prerequisite for capturing the value that the technology can deliver.

From the Author

I didn’t come to AI leadership through a boardroom. I came to it through the exam room.

Building healthcare apps taught me something that no MBA course ever did: the distance between a good idea and a tool that works in a clinical environment at 3 a.m. is measured in a thousand small decisions that have nothing to do with the algorithm and everything to do with the humans using it. Font size matters when a nurse is reading it under fluorescent lights after a 12-hour shift. Load time matters when a physician has 90 seconds between patients. Language matters when your patient speaks Spanish and your interface doesn’t.

I built AI-powered discharge instructions that translate into 54 languages. I created coloring book discharge summaries for pediatric patients, sitting with kids and colored alongside them to teach asthma management, because a four-year-old doesn’t read a medication guide, but a four-year-old will remember what they colored. I developed sign language video instructions for patients who are hard of hearing. None of this was theoretical. All of it was built, tested, revised, and deployed.

Leading teams through AI adoption requires the same skill set: specificity, patience, and the willingness to fail in front of the people you’re asking to change. When I tell a clinical team that AI will augment their practice, I’m not reading from a vendor deck. I’m speaking as someone who built the tools, used the tools, watched the tools break, fixed them, and deployed them again. That credibility, earned in the wreckage of Version 1.0, is the only currency that matters when you’re asking a skeptical physician to trust a machine.

Key Takeaways

Clinical leadership is the decisive variable in healthcare AI success. Clinician-led AI deployments have 7.79 times higher odds of significant impact compared to technologist-led deployments. Every

CEO must ensure CMO co-leadership in AI governance from day one [18].

Governance infrastructure is not optional: it is a compliance necessity. The 66 percent of hospitals with an AI committee still leave one-third of hospitals without one. In an environment with more than 1,250 FDA-authorized AI devices and multiple new federal transparency requirements, a formal governance structure is the minimum responsible standard, not a best practice aspiration [8].

The three-role framework (CEO, CIO, CMO) creates clear accountability and prevents the most common structural failure in health system AI programs. The CEO owns vision and board reporting; the CIO owns infrastructure, security, and data governance; the CMO owns clinical validation, safety monitoring, and physician engagement. Overlapping responsibilities without clear ownership is where programs collapse.

Vendor validation claims require independent verification before any clinical deployment. The Epic sepsis model, Optum bias case, and IBM Watson failures each resulted from accepting vendor performance claims at face value. Every deployment must include local validation on the deploying institution's patient population before entering a clinical workflow.

Cybersecurity is inseparable from AI governance. The Change Healthcare attack, which cost $2.87 billion and exposed 190 million patient records through a single missing multifactor authentication control, demonstrates that AI systems are nodes in a shared digital infrastructure, not isolated tools [30].

ROI measurement must be built into AI programs at the outset, not added after the fact. Radiology AI has demonstrated 451 percent to 791 percent five-year returns in peer-reviewed research [61]. Ambient AI scribes have produced documented burnout reductions with multimillion-dollar financial implications [63]. These figures are achievable only when organizations define KPIs, collect baselines, and measure systematically.

Revenue cycle AI represents the highest near-term financial opportunity for most health systems. McKinsey estimates $60 million to $120 million in annual savings per large health system from revenue cycle AI, deployed against a $140 billion annual industry cost base [65]. The organizations that capture these gains will fund the governance infrastructure to extend AI responsibly across clinical domains.

State and federal regulatory requirements are accelerating and must be integrated into governance structures now. California SB 1120, Colorado SB 24-205, ONC HTI-1, and the EU AI Act's extraterritorial reach are all active or imminent. Legal and compliance leaders must be integrated into AI governance committees, not consulted after deployment decisions are made [41–43].

The chief AI officer role is becoming standard infrastructure, not a luxury. Health systems appointing dedicated executive AI leadership are building a structural competitive advantage. Those that do not risk the fragmented, uncoordinated adoption that peer-reviewed research identifies as the most dangerous organizational failure mode.

The window for deliberate governance is narrowing as the AI divide widens. The Deloitte divide is accelerating. Organizations that act now, with clear governance, validated tools, and executive accountability, will compound their advantages over the next three to five years, while organizations still deliberating will be competing against health systems already two or more years into production operation.

References

1. Grand View Research. AI in healthcare market size, share & trends analysis report [Internet]. San Francisco (CA): Grand View Research; 2024 [cited 2026 Jul 4]. Available from: https://www.grandviewresearch.com/press-release/global-artificial-intelligence-healthcare-market.

2. Chang W, Owusu-Mensah P, Everson J, Richwine C. Hospital trends in the use, evaluation, and governance of predictive AI, 2023-2024.

ASTP/ONC Data Brief No. 80 [Internet]. Washington (DC): Office of the Assistant Secretary for Technology Policy; 2025 Sep [cited 2026 Jul 4]. Available from: https://www.healthit.gov/data/data-briefs/hospital-trends-use-evaluation-and-governance-predictive-ai-2023-2024.

3. Deloitte. Deloitte report reveals an emerging 'AI divide' in health care [Internet]. New York: PR Newswire; 2026 Feb 11 [cited 2026 Jul 4]. Available from: https://www.prnewswire.com/news-releases/deloitte-report-reveals-an-emerging-ai-divide-in-health-care-302684719.html.

4. American Medical Association. Physician enthusiasm grows for health care AI [Internet]. Chicago (IL): American Medical Association; 2025 Feb 12 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/press-center/ama-press-releases/ama-physician-enthusiasm-grows-health-care-ai.

5. Silicon Valley Bank. AI investment accounted for nearly half of healthcare investment in 2025; Silicon Valley Bank releases 17th Healthcare Investments and Exits Report [Internet]. Santa Clara (CA): PR Newswire; 2026 Jan 8 [cited 2026 Jul 4]. Available from: https://www.prnewswire.com/news-releases/ai-investment-accounted-for-nearly-half-of-healthcare-investment-in-2025-silicon-valley-bank-releases-17th-healthcare-investments-and-exits-report-302656179.html.

6. Bipartisan Policy Center. FDA oversight: understanding the regulation of health AI tools [Internet]. Washington (DC): Bipartisan Policy Center; 2025 [cited 2026 Jul 4]. Available from: https://bipartisanpolicy.org/issue-brief/fda-oversight-understanding-the-regulation-of-health-ai-tools.

7. Chang W, Owusu-Mensah P, Everson J, Richwine C. Hospital trends in the use, evaluation, and governance of predictive AI, 2023-2024. ASTP/ONC Data Brief No. 80 [Internet]. Washington (DC): Office of the Assistant Secretary for Technology Policy; 2025 Sep [cited 2026 Jul 4]. Available from: https://www.healthit.gov/data/data-

briefs/hospital-trends-use-evaluation-and-governance-predictive-ai-2023-2024.

8. American Hospital Association. AHA response to HHS RFI on AI in health care [Internet]. Chicago (IL): American Hospital Association; 2026 Feb 23 [cited 2026 Jul 4]. Available from: https://www.aha.org/lettercomment/2026-02-23-aha-response-hhs-rfi-ai-health-care.

9. Paubox. What is the AMA AI governance framework? [Internet]. Houston (TX): Paubox; 2025 [cited 2026 Jul 4]. Available from: https://www.paubox.com/blog/what-is-the-ama-ai-governance-framework.

10. McKinsey & Company. The coming evolution of healthcare AI toward a modular architecture [Internet]. New York: McKinsey & Company; 2025 [cited 2026 Jul 4]. Available from: https://www.mckinsey.com/industries/healthcare/our-insights/the-coming-evolution-of-healthcare-ai-toward-a-modular-architecture.

11. Digital Health Insights. McKinsey's 2025 tech trends report finds healthcare caught between AI promise and perils [Internet]. Digital Health Insights; 2025 [cited 2026 Jul 4]. Available from: https://dhinsights.org/blog/mckinseys-2025-tech-trends-report-finds-healthcare-caught-between-ai-promise-and-perils.

12. Walker L. BCG's CEO roadmap for agentic AI [Internet]. Enterprise AI Executive; 2025 Dec 21 [cited 2026 Jul 4]. Available from: https://enterpriseaiexecutive.ai/p/bcg-unveils-ceo-roadmap-for-agentic-ai.

13. Deloitte. Generative AI to reshape the future of health care [Internet]. Deloitte; 2025 [cited 2026 Jul 4]. Available from: https://www.deloitte.com/us/en/Industries/life-sciences-health-care/articles/generative-ai-in-healthcare.html.

14. Vention. AI in healthcare statistics [Internet]. Vention; 2025 [cited 2026 Jul 4]. Available from: https://ventionteams.com/healthtech/ai/statistics.

15. PrivaPlan. Healthcare leads in AI adoption: new reports show [Internet]. PrivaPlan; 2025 [cited 2026 Jul 4]. Available from: https://privaplan.com/healthcare-leads-in-ai-adoption-new-reports-show.

16. Menlo Ventures. 2025: the state of AI in healthcare [Internet]. San Francisco (CA): Menlo Ventures; 2025 Oct 21 [cited 2026 Jul 4]. Available from: https://menlovc.com/perspective/2025-the-state-of-ai-in-healthcare.

17. U.S. Food and Drug Administration. Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions [Internet]. Silver Spring (MD): FDA; 2024 Dec [cited 2026 Jul 4]. Available from: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial-intelligence.

18. Li Q, Li P, Hao H, Padman R, Gao H, Travis Gossey J, et al. The impact of leadership on AI deployment study outcomes in healthcare: an integrative analysis. npj Digit Med. 2025 Nov 28;8(1):799. doi:10.1038/s41746-025-02177-x. Available from: https://www.nature.com/articles/s41746-025-02177-x.

19. Healthcare IT News. Kaiser Permanente's new head of AI on 'two fundamental shifts' the technology will enable [Internet]. Healthcare IT News; 2025 Jul 9 [cited 2026 Jul 4]. Available from: https://www.healthcareitnews.com/news/kaiser-permanentes-new-head-ai-two-fundamental-shifts-technology-will-enable.

20. Kaiser Permanente. AI in health care: 7 principles of responsible use [Internet]. Oakland (CA): Kaiser Permanente; 2025 [cited 2026 Jul 4]. Available from: https://about.kaiserpermanente.org/news/ai-in-health-care-7-principles-of-responsible-use.

21. Cleveland Clinic. Cleveland Clinic names first chief AI officer [Internet]. Cleveland (OH): Cleveland Clinic Newsroom; 2024 Jul 29 [cited 2026 Jul 4]. Available from:

https://newsroom.clevelandclinic.org/2024/07/29/cleveland-clinic-names-first-chief-ai-officer.

22. Cleveland Clinic. Less typing, more talking: how ambient AI is reshaping clinical workflow at Cleveland Clinic [Internet]. Cleveland (OH): Consult QD, Cleveland Clinic; 2025 [cited 2026 Jul 4]. Available from: https://consultqd.clevelandclinic.org/less-typing-more-talking-how-ambient-ai-is-reshaping-clinical-workflow-at-cleveland-clinic.

23. Bean R. How Cleveland Clinic is innovating in healthcare with data, analytics, and AI. Forbes [Internet]. 2024 Oct 6 [cited 2026 Jul 4]. Available from: https://www.forbes.com/sites/randybean/2024/10/6/how-cleveland-clinic-is-innovating-in-healthcare-with-data-analytics-and-ai.

24. Cleveland Clinic. Cleveland Clinic and G42 to advance healthcare through artificial intelligence [Internet]. Cleveland (OH): Cleveland Clinic Newsroom; 2025 Mar 25 [cited 2026 Jul 4]. Available from: https://newsroom.clevelandclinic.org/2025/03/25/cleveland-clinic-and-g42-to-advance-healthcare-through-artificial-intelligence.

25. NYC Health + Hospitals. NYC Health + Hospitals announces two key leadership appointments [Internet]. New York: NYC Health + Hospitals; 2025 [cited 2026 Jul 4]. Available from: https://www.nychealthandhospitals.org/pressrelease/nyc-health-hospitals-announces-two-key-leadership-appointments.

26. Washington University School of Medicine. WashU Medicine, BJC Health System launch Center for Health AI [Internet]. St. Louis (MO): WashU Medicine; 2024-2025 [cited 2026 Jul 4]. Available from: https://medicine.washu.edu/news/washu-medicine-bjc-health-system-launch-center-for-health-ai.

27. Dolfing H. Case study 20: the $4 billion AI failure of IBM Watson for Oncology [Internet]. Henrico Dolfing; 2025 [cited 2026 Jul 4]. Available

from: https://www.henricodolfing.ch/case-study-20-the-4-billion-ai-failure-of-ibm-watson-for-oncology.

28. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019 Oct 25; 366(6464):447-53. doi:10.1126/science.aax2342. Available from: https://pubmed.ncbi.nlm.nih.gov/31649194.

29. Wong A, Otles E, Donnelly JP, Krumm A, McCullough J, DeTroyer-Cooley O, et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Intern Med. 2021 Aug 1; 181(8):1065-70. doi:10.1001/jamainternmed.2021.2626. Available from: https://pubmed.ncbi.nlm.nih.gov/34152373.

30. HIPAA Journal. Change Healthcare responding to cyberattack [Internet]. HIPAA Journal; 2025 [cited 2026 Jul 4]. Available from: https://www.hipaajournal.com/change-healthcare-responding-to-cyberattack.

31. American Hospital Association. Change Healthcare cyberattack underscores urgent need to strengthen cyber preparedness for individual health care organizations and the entire health care system [Internet]. Chicago (IL): American Hospital Association; 2024 [cited 2026 Jul 4]. Available from: https://www.aha.org/change-healthcare-cyberattack-underscores-urgent-need-strengthen-cyber-preparedness-individual-health-care-organizations-and.

32. AndChange. Change readiness in healthcare's digital transformation [Internet]. AndChange; 2025 [cited 2026 Jul 4]. Available from: https://www.andchange.com/change-readiness-healthcare-digital-transformation.

33. HealthTech Magazine. Upskilling healthcare IT staff in AI and cybersecurity [Internet]. HealthTech Magazine; 2025 May [cited 2026 Jul 4]. Available from: https://healthtechmagazine.net/article/2025/05/upskilling-healthcare-it-staff-ai-cybersecurity-perfcon.

34. Fast Company. AI's impact: transforming the build, buy, or partner decision framework in healthcare [Internet]. Fast Company; 2025 [cited

2026 Jul 4]. Available from: https://www.fastcompany.com/91324807/ais-impact-transforming-the-build-buy-or-partner-decision-framework-in-healthcare.

35. U.S. Food and Drug Administration, Health Canada, Medicines and Healthcare products Regulatory Agency. Good machine learning practice for medical device development: guiding principles [Internet]. Silver Spring (MD): FDA; 2021 [cited 2026 Jul 4]. Available from: https://www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles.

36. Bipartisan Policy Center. FDA oversight: understanding the regulation of health AI tools [Internet]. Washington (DC): Bipartisan Policy Center; 2025 [cited 2026 Jul 4]. Available from: https://bipartisanpolicy.org/issue-brief/fda-oversight-understanding-the-regulation-of-health-ai-tools.

37. Office of the National Coordinator for Health IT. Health data, technology, and interoperability: certification program updates, algorithm transparency, and information sharing (HTI-1) final rule [Internet]. Washington (DC): ONC; 2024 [cited 2026 Jul 4]. Available from: https://www.healthit.gov/regulations/hti-rules/hti-1-final-rule.

38. Bipartisan Policy Center. Paying for AI in U.S. health care [Internet]. Washington (DC): Bipartisan Policy Center; 2026 [cited 2026 Jul 4]. Available from: https://bipartisanpolicy.org/issue-brief/paying-for-ai-in-u-s-health-care.

39. STAT News. Medicare picks tech firms to run AI prior authorization pilot [Internet]. STAT; 2025 Nov 6 [cited 2026 Jul 4]. Available from: https://www.statnews.com/2025/11/06/medicare-wiser-prior-authorization-pilot-tech-vendors.

40. American Medical Association. New voluntary CMS pay model encourages use of health tech [Internet]. Chicago (IL): American Medical Association; 2025 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/payment-

delivery-models/new-voluntary-cms-pay-model-encourages-use-health-tech.

41. Fenwick & West LLP. California's SB 1120 regulates AI in health plan utilization review and management activities starting in January [Internet]. Fenwick & West LLP; 2024 [cited 2026 Jul 4]. Available from: https://www.fenwick.com/insights/publications/californias-sb-1120-regulates-ai-in-health-plan-utilization-review-and-management-activities-starting-in-january.

42. Colorado General Assembly. SB24-205: consumer protections for artificial intelligence [Internet]. Denver (CO): Colorado General Assembly; 2024 [cited 2026 Jul 4]. Available from: https://leg.colorado.gov/bills/sb24-205.

43. van Kolfschooten H, van Oirschot J. The EU Artificial Intelligence Act (2024): implications for healthcare. Health Policy. 2024; 149:105152. doi:10.1016/j.healthpol.2024.105152. Available from: https://pubmed.ncbi.nlm.nih.gov/39244818.

44. Harvard Medical School Center for Bioethics. From code to conscience: an ethical framework for healthcare AI [Internet]. Boston (MA): Harvard Medical School Center for Bioethics; 2025 Nov 13 [cited 2026 Jul 4]. Available from: https://www.ethics.harvard.edu/news/2025/11/code-conscience-ethical-framework-healthcare-ai-0.

45. World Health Organization. WHO issues first global report on AI in health and six guiding principles for its design and use [Internet]. Geneva: World Health Organization; 2021 Jun 28 [cited 2026 Jul 4]. Available 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.

46. Organisation for Economic Co-operation and Development. Recommendation of the Council on artificial intelligence (OECD AI Principles) [Internet]. Paris: OECD; 2019 [cited 2026 Jul 4]. Available from: https://www.oecd.org/en/topics/ai-principles.html.

47. White House Office of Science and Technology Policy. Blueprint for an AI Bill of Rights [Internet]. Washington (DC): The White House; 2022 [cited 2026 Jul 4]. Available from: https://www.whitehouse.gov/ostp/ai-bill-of-rights.

48. American Medical Association. AMA issues new principles for AI development, deployment & use [Internet]. Chicago (IL): American Medical Association; 2023 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/press-center/ama-press-releases/ama-issues-new-principles-ai-development-deployment-use.

49. American Medical Association. AMA STEPS Forward: governance for augmented intelligence [Internet]. Chicago (IL): American Medical Association; 2025 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/sustainability/ama-steps-forward-governance-augmented-intelligence.

50. ECRI. Artificial intelligence tops 2025 health technology hazards list [Internet]. Plymouth Meeting (PA): ECRI; 2024 [cited 2026 Jul 4]. Available from: https://home.ecri.org/blogs/ecri-news/artificial-intelligence-tops-2025-health-technology-hazards-list.

51. ECRI. Misuse of AI chatbots tops annual list of health technology hazards [Internet]. Plymouth Meeting (PA): ECRI; 2026 [cited 2026 Jul 4]. Available from: https://home.ecri.org/blogs/ecri-news/misuse-of-ai-chatbots-tops-annual-list-of-health-technology-hazards.

52. American Hospital Association Center for Health Innovation. The C-suite rewrites the org chart: how health systems are elevating technology [Internet]. Chicago (IL): AHA Center for Health Innovation Market Scan; 2026 Feb 3 [cited 2026 Jul 4]. Available from: https://www.aha.org/aha-center-health-innovation-market-scan/2026-02-03-c-suite-rewrites-org-chart-how-health-systems-are-elevating-technology.

53. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1 [Internet].

Gaithersburg (MD): NIST; 2023 Jan [cited 2026 Jul 4]. doi:10.6028/NIST.AI.100-1. Available from: https://www.nist.gov/itl/ai-risk-management-framework.

54. College of Healthcare Information Management Executives (CHIME). CHIME releases ten principles for AI in healthcare [Internet]. Ann Arbor (MI): CHIME; 2024 [cited 2026 Jul 4]. Available from: https://chimecentral.org/chime-releases-ten-principles-for-ai-in-healthcare.

55. Li Q, Li P, Hao H, Padman R, Gao H, Travis Gossey J, et al. The impact of leadership on AI deployment study outcomes in healthcare: an integrative analysis. npj Digit Med. 2025 Nov 28;8(1):799. doi:10.1038/s41746-025-02177-x. Available from: https://www.nature.com/articles/s41746-025-02177-x.

56. FTI Consulting. Overcoming AI misconceptions: the role of governance in healthcare [Internet]. FTI Consulting; 2024 [cited 2026 Jul 4]. Available from: https://www.fticonsulting.com/insights/articles/overcoming-ai-misconceptions-role-governance-healthcare.

57. Kim JY, Hasan A, Kueper J, Tang T, Hayes C, Fine B, et al. Establishing organizational AI governance in healthcare. npj Digit Med. 2025; 8:522. doi:10.1038/s41746-025-01909-3. PMCID: PMC12356831.

58. Public Health AI Handbook. Vendor AI evaluation toolkit [Internet]. Public Health AI Handbook; 2026 [cited 2026 Jul 4]. Available from: https://publichealthaihandbook.com/appendices/vendor-evaluation.html.

59. Digital Medicine Society (DiMe). Health AI maturity model [Internet]. Boston (MA): Digital Medicine Society; 2025 [cited 2026 Jul 4]. Available from: https://dimesociety.org/ai-implementation-in-healthcare-playbook/ai-evaluation-readiness/health-ai-maturity-model.

60. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK; SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the

CONSORT-AI extension. BMJ. 2020; 370:m3164. doi:10.1136/bmj.m3164. Available from: https://www.bmj.com/content/370/bmj.m3164.

61. Bharadwaj P, Nicola L, Breau-Brunel M, Sensini F, Tanova-Yotova N, Atanasov P, et al. Unlocking the value: quantifying the return on investment of hospital artificial intelligence. J Am Coll Radiol. 2024 Oct; 21(10):1677-85. doi:10.1016/j.jacr.2024.02.034. Available from: https://pubmed.ncbi.nlm.nih.gov/38499053.

62. American Medical Association. AI scribes save 15,000 hours—and restore the human side of medicine [Internet]. Chicago (IL): American Medical Association; 2025 Jun 12 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/ai-scribes-save-15000-hours-and-restore-human-side-medicine.

63. Mass General Brigham. Ambient documentation technologies reduce physician burnout [Internet]. Boston (MA): Mass General Brigham Newsroom; 2025 Aug 21 [cited 2026 Jul 4]. Available from: https://www.massgeneralbrigham.org/en/about/newsroom/press-releases/ambient-documentation-technologies-reduce-physician-burnout.

64. Yale School of Medicine. AI scribes reduce physician burnout, return focus to the patient [Internet]. New Haven (CT): Yale School of Medicine; 2025 [cited 2026 Jul 4]. Available from: https://medicine.yale.edu/news-article/ai-scribes-reduce-physician-burnout-return-focus-to-the-patient.

65. McKinsey & Company. Agentic AI and the race to a touchless revenue cycle [Internet]. New York: McKinsey & Company; 2026 Jan 9 [cited 2026 Jul 4]. Available from: https://www.mckinsey.com/industries/healthcare/our-insights/agentic-ai-and-the-race-to-a-touchless-revenue-cycle.

66. AKASA. Survey: cost-to-collect nearly 0.25% lower for hospitals using automation [Internet]. South San Francisco (CA): AKASA; 2022 Sep 15

[cited 2026 Jul 4]. Available from: https://akasa.com/press/survey-cost-to-collect-lower-with-automation.

67. HealthManagement.org. The return on investment of artificial intelligence in radiology: is it really profitable? HealthManagement [Internet]. 2026 [cited 2026 Jul 4]. Available from: https://healthmanagement.org/c/healthmanagement/issuearticle/the-return-on-investment-of-artificial-intelligence-in-radiology-is-it-really-profitable.

68. Grand View Research. U.S. artificial intelligence in healthcare market report [Internet]. San Francisco (CA): Grand View Research; 2024 [cited 2026 Jul 4]. Available from: https://www.grandviewresearch.com/industry-analysis/us-artificial-intelligence-healthcare-market-report.

69. MarketsandMarkets. AI in hospital operations market worth $7.51 billion in 2025 [Internet]. MarketsandMarkets; 2025 [cited 2026 Jul 4]. Available from: https://www.marketsandmarkets.com/PressReleases/ai-in-hospital-operations.asp.

70. Azevedo MA, Rider J. Crunchbase sector snapshot: funding to AI-related healthcare companies rises in 2025 [Internet]. Crunchbase News; 2025 Nov 10 [cited 2026 Jul 4]. Available from: https://news.crunchbase.com/health-wellness-biotech/ai-healthcare-funding-rises-2025.

71. American Hospital Association Center for Health Innovation. The C-suite rewrites the org chart: how health systems are elevating technology [Internet]. Chicago (IL): AHA Center for Health Innovation Market Scan; 2026 Feb 3 [cited 2026 Jul 4]. Available from: https://www.aha.org/aha-center-health-innovation-market-scan/2026-02-03-c-suite-rewrites-org-chart-how-health-systems-are-elevating-technology.

72. Deloitte. AI could help health plans simplify prior authorization and comply with regulations [Internet]. Deloitte; 2026 [cited 2026 Jul 4]. Available from: https://www.deloitte.com/us/en/Industries/life-

sciences-health-care/blogs/health-care/ai-could-help-health-plans-simplify-prior-auth-comply-with-regs.html.

73. van Kolfschooten H, van Oirschot J. The EU Artificial Intelligence Act (2024): implications for healthcare — Article 99 penalties [Internet]. Artificial Intelligence Act EU / EU AI Act text; 2024 [cited 2026 Jul 4]. Available from: https://artificialintelligenceact.eu/article/99.

74. You JG, Dbouk RH, Landman A, Ting DY, Dutta S, Wang JC, et al. Ambient documentation technology in clinician experience of documentation burden and burnout. JAMA Netw Open. 2025; 8(8):e2528056. doi:10.1001/jamanetworkopen.2025.28056. PMID: 40839265.

AI in Emergency Medicine: What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution, First Edition. Harvey Castro.
© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

10
Ethics, Equity, and AI Accountability

My mother was 13 when she arrived in this country from Colombia. She didn’t speak the language. She didn’t have a high school diploma. She had me when she was 16.

Every system we encountered, from the welfare office to the public school to the emergency room where a doctor stitched up my hand without pain medication and never looked me in the eye, treated us like data points to be processed, not people to be cared for. We weren’t seen as individuals with specific needs. We were a demographic category. Low-income. Minority. Uninsured. The system gave us what the system was built to give people in our category, and it wasn’t much.

When I read the research on algorithmic bias in healthcare, the Obermeyer study showing that a commercial algorithm used cost as a proxy for illness and undertreated Black patients at massive scale, or the pulse oximetry data showing that Black patients received dangerously inaccurate oxygen readings that delayed treatment, I don’t experience it as an abstract policy problem. I experience it as the digitization of what my mother and I lived through. An algorithm that uses cost to predict health need will conclude that people who can’t afford care are healthier than people who can. That’s not a bug in the math. That’s the math encoding the same structural inequality my mother faced when she brought a bleeding child to an emergency room and waited eight hours because no one thought we were worth seeing faster.

That’s why this chapter exists, and that’s why accountability in AI is not a bureaucratic exercise for me. It’s personal. Either the tools we build will replicate the system that made my mother invisible or they’ll be built differently, with equity as a design requirement, not an afterthought. The evidence in this chapter shows that both outcomes are possible. Recalibrating the Optum algorithm eliminated the racial enrollment gap. Fine-tuning AI dermatology tools on diverse images erased skin-tone performance disparities. The tools can be fair. The question is whether the

people who deploy them choose to make them fair and whether governance frameworks exist to hold them accountable when they don’t.

Executive Summary

Artificial intelligence is reshaping healthcare at an unprecedented pace and scale. More than 1,250 AI-enabled medical devices have been authorized by the FDA [1] and a 2024 American Medical Association survey found that 66 percent of physicians used AI systems in clinical practice, nearly double the 38 percent recorded just one year earlier [2]. The technology holds genuine promise: earlier disease detection, reduced burnout, and care delivered where no specialist has ever set foot. But promise and peril travel together.

In 2019, a landmark study published in *Science* revealed that a widely used commercial risk-prediction algorithm (later identified as Optum’s Impact Pro tool, which directly guided care for more than 70 million Americans and belonged to a broader class of similar tools applied to roughly 200 million Americans per year) was systematically undertreating Black patients [3]. The algorithm used healthcare costs as a proxy for health need. Because systemic inequities cause Black patients to spend less on healthcare than equally sick white patients, the model concluded Black patients were healthier than they were, cutting their enrollment in high-risk care management programs from a corrected figure of 46.5 percent down to 17.7 percent [3]. The algorithm did not use race as an input variable. The bias was invisible until researchers looked. That is the central problem this chapter addresses.

The ethical challenges of healthcare AI are not abstractions for a future reckoning. They are present-day realities. ECRI ranked “Risks with AI-enabled health technologies” as the number-one health technology hazard for 2025, the first time artificial intelligence has topped that list [4]. The Federation of State Medical Boards declared in May 2024 that physicians bear ultimate responsibility for any AI-related clinical error [5]. A Senate investigation documented a ninefold increase in post-acute care denials at one major insurer coinciding with the adoption of an AI prediction tool [6]. A landmark colonoscopy study found that after physicians worked with AI, their unassisted detection skills measurably declined [7].

This chapter examines the evidence on algorithmic bias, documents the key ethical case studies, maps the regulatory field, and offers a practical leadership framework for building health AI systems that are effective, fair, transparent, and accountable. The goal is not to slow AI adoption. It's to ensure that the revolution in care benefits every patient, not just those who happen to resemble the training data.

Current Evidence

The deployment of AI across healthcare has reached a scale that makes the ethical questions in this chapter urgent rather than theoretical. Before examining those ethical dimensions, it's worth understanding the scope of what's already in clinical use and who it affects.

The Scale of Deployment

The numbers alone demand ethical attention. The FDA's public database lists more than 1,250 AI-enabled medical devices authorized for marketing in the United States as of July 2025, up from approximately 950 devices just 12 months earlier [1]. A peer-reviewed systematic review in *JAMA Network Open* found that of the first 950 authorized AI/ML devices, 97 percent were cleared via the 510(k) pathway, a mechanism that does not require independent clinical data demonstrating safety in the same way that full premarket approval does [8]. Radiology dominates the portfolio, accounting for 76 percent of approvals, yet the applications reach into every specialty: pathology, cardiology, ophthalmology, and emergency medicine [8].

Hospitals are adopting fast. A 2024 American Hospital Association survey found that 71 percent of hospitals were using predictive AI tools of some kind [9]. The AMA's survey found that 66 percent of physicians used AI systems in 2024 [2] and the AMA released a formal eight-step governance toolkit in August 2025, acknowledging that "healthcare AI technology is evolving faster than hospitals can implement tools" [10]. Against this pace of adoption, ECRI's declaration that AI is the top health technology hazard for 2025 is not alarmism. It's calibration [4].

Documented Bias: Three Landmark Cases

The three cases that follow represent the most thoroughly documented instances of algorithmic bias in healthcare. Each reveals a different mechanism by which AI systems can produce discriminatory outcomes, and each has generated regulatory, legal, and institutional responses that continue to shape the field.

Optum and the Cost-as-Proxy Problem The Obermeyer et al. study in *Science* (2019) examined 43,539 white patients and 6,079 Black patients enrolled in risk-based contracts at a large academic hospital [3]. At any given algorithmic risk score, Black patients were considerably sicker than white patients: those at the 97th percentile threshold had 26.3 percent more chronic conditions than equally scored white patients (4.8 versus 3.8 distinct conditions; $P < 0.001$ [3]. The root cause was structural: the algorithm predicted future healthcare costs, not illness severity. Because Black patients receive less care than equally sick white patients, owing to transportation barriers, lower healthcare trust, reduced access, and the cumulative effects of racial discrimination, lower predicted costs signaled falsely lower health need. The algorithm did not use race as a variable, yet it encoded racial inequity through every dollar in its training data [3]. The study estimated this class of commercial risk-prediction tools affects roughly 200 million Americans annually [3].

Dermatology AI and the Skin-Tone Gap Daneshjou et al. published in *Science Advances* (2022) the first study using a pathologically confirmed, diverse skin-tone dataset to benchmark dermatology AI [11]. The findings were stark. DeepDerm's sensitivity for detecting malignancy was 0.69 on light skin (Fitzpatrick types I–II) versus 0.23 on dark skin (Fitzpatrick types V–VI), a nearly threefold disparity (Fisher's exact $P = 5.65 \times 10^{-6}$) [11]. ModelDerm showed 0.41 versus 0.12 sensitivity across the same skin-tone groups [11]. The underlying cause was training data: in the widely used ISIC dataset, more than 70 percent of images depict light skin while fewer than 8 percent represent dark skin [12]. These gaps are not just technical. Between 2011 and 2015, the five-year melanoma survival rate was 66 percent for Black patients compared to 90 percent for non-Hispanic white patients, a disparity that AI performing worse on dark skin can widen further [13]. The Daneshjou study also demonstrated the path to remedy:

fine-tuning DeepDerm on the diverse dataset caused its performance on dark skin to actually exceed its performance on light skin, with Roc-AUC of 0.76 versus 0.73 [11]. I've seen this pattern in my own work: diverse training data isn't a bonus feature. It's the difference between a tool that works and one that harms.

Pulse Oximetry and the Compounding of Hardware Bias A December 2020 study in the *New England Journal of Medicine* by Sjoding et al. found that Black patients had nearly three times the frequency of occult hypoxemia (true oxygen saturation below 88 percent despite pulse oximeter readings of 92–96 percent) compared to white patients: 11.7 percent versus 3.6 percent in the University of Michigan cohort [14]. A follow-up multicenter study of 128,285 paired measurements found Black patients had 65 percent higher adjusted odds of occult hypoxemia (OR 1.65; 95% CI 1.28–2.14; $P < 0.001$) [15]. The mechanism is melanin's effect on light absorption at clinically critical oxygen saturation ranges. The direct relevance to AI is structural: sepsis prediction algorithms, deterioration alert systems, and oxygen therapy decision tools that incorporate pulse oximetry data inherit this racial inaccuracy at scale, potentially compounding disparities across every hospital department using AI-driven deterioration monitoring [16].

Ethics Frameworks: From Principles to Practice Healthcare

AI ethics draws from two principal global frameworks and one foundational bioethics tradition. In June 2021, the World Health Organization published "Ethics and Governance of Artificial Intelligence for Health," the first global guidance document produced after 18 months of consultation with international experts [17]. WHO established six core principles:

- Protecting human autonomy
- Promoting human well-being, safety, and the public interest
- Ensuring transparency, explainability, and intelligibility
- Building responsibility and accountability
- Ensuring inclusiveness and equity

- Promoting AI that is responsive and sustainable [17]

Principle 5, inclusiveness and equity, explicitly requires that AI for health be designed for equitable use regardless of age, sex, gender, income, race, ethnicity, sexual orientation, or ability [17].

UNESCO's 193 member states unanimously adopted the "Recommendation on the Ethics of Artificial Intelligence" in November 2021, the world's first global standard-setting instrument on AI ethics [18]. Its 10 guiding principles include proportionality and do no harm, fairness and nondiscrimination, transparency and explainability, responsibility and accountability, and multistakeholder governance [18]. On healthcare specifically, UNESCO calls on states to ensure that final decisions are made by humans, require patient informed consent for AI-assisted care, invest in ethics review committees, and develop guidelines for human-AI interactions [18].

Beauchamp and Childress's four principles of biomedical ethics (autonomy, beneficence, nonmaleficence, and justice) give the clinical translation layer [19]. Applied to AI:

- Autonomy requires that patients be informed when AI influences their clinical decisions.
- Beneficence asks whether AI tools demonstrably improve outcomes for all populations served.
- Nonmaleficence addresses the harms of algorithmic errors and automation bias.
- Justice confronts algorithmic disparities and the digital divide.

A 2025 analysis in "Voices in Bioethics" (Columbia University) applying this framework to generative AI found specific risks to all four principles, concluding that without appropriate governance, generative AI integration can lead to data protection violations, automation bias, unclear accountability, healthcare inequality, and incorrect clinical decisions [20].

The Urgency Signal

The AMA's near-doubling of physician AI adoption from 38 percent to 66 percent in a single year [2] represents the fastest documented diffusion of a clinical technology in the AMA's survey history. ECRI's ranking of AI as the top health technology hazard for 2025, and "insufficient governance of AI in healthcare" as the second-highest patient safety concern, confirms that adoption has outpaced the ethical and safety infrastructure needed to manage it [4]. The imperative is not to pause but to build.

Real-World Case Studies

Three cases define the ethical failure modes that healthcare leaders must understand and prevent. Each reveals a distinct dimension of the problem.

Case 1: Optum (When Bias Wears the Face of Neutrality)

The Optum case, which we discussed earlier in this book in the context of AI governance failures, is the canonical example of algorithmic bias in healthcare and deserves close examination through an ethical lens. The Optum case is the canonical example of a race-neutral algorithm that nonetheless encoded racial discrimination at population scale. The mechanism, as documented by Obermeyer et al. in *Science*, was conceptually simple and practically catastrophic: costs were used as a proxy for health need [3]. The algorithm predicted how much a patient would cost the health system in the coming year and used that prediction to decide who received intensive care management. The assumption was that sicker patients cost more. For white patients, that assumption held reasonably well. For Black patients, it did not, because decades of systemic inequity meant that equally sick Black patients received less care and therefore incurred lower costs [3].

The numbers were unambiguous. At the high-risk threshold automatically triggering care management enrollment, Black patients had 26.3 percent more chronic conditions than equally scored white patients [3]. Correcting the bias by substituting illness severity (number of active chronic conditions) for cost as the outcome variable would have more than doubled Black patient enrollment in care management programs, from 17.7 percent

to 46.5 percent [3]. The algorithm affected an estimated 200 million people annually [3].

The Optum case is instructive not just because bias was found but because of how it was hidden and then addressed. The bias was invisible because it arose from a training variable, healthcare cost, that appeared objective. No one reviewing the algorithm's inputs would have flagged "healthcare expenditures" as a racial variable; it's a standard administrative data field. The bias was uncovered only because researchers from UC Berkeley and Harvard analyzed outcomes by race and found the disparity [3]. When Obermeyer shared findings with Optum, the company acknowledged the issue and stated it was working on a revised algorithm [21]. Even well-intentioned, sophisticated developers can embed systemic inequity into a widely deployed tool. Independent scrutiny, not vendor assurance, is the essential safeguard.

The Optum case also established a replicable remediation pathway. By substituting direct health measures for proxy variables, by expanding the definition of what the algorithm is trying to predict, the racial performance gap was substantially eliminated [3]. The lesson for health system leaders: the choice of outcome variable is an ethical decision, not just a technical one.

Case 2: DeepMind and Royal Free Hospital (When Data Partnership Becomes Data Extraction)

In 2015, the Royal Free London NHS Foundation Trust shared the identifiable personal health records of approximately 1.6 million patients with Google's DeepMind subsidiary, without patients' knowledge or explicit consent, for the purpose of testing an app called Streams designed to detect acute kidney injury [22]. The data included not only kidney-relevant clinical information but also HIV status, drug overdose records, and abortion records [22]. Patients had not been informed their data was being used for this purpose, and the vast majority had never been evaluated or treated for kidney disease, meaning the data transfer could not be justified even under the broadly interpreted "direct care" implied consent doctrine that the Trust initially relied upon [23].

On July 3, 2017, the UK Information Commissioner's Office ruled that the Royal Free NHS Foundation Trust had failed to comply with the Data Protection Act 1998 [24]. The Information Commissioner's Office (ICO, the United Kingdom's independent data protection authority) found that patients had not been adequately informed their data was being used; that the processing was neither fair nor transparent; that the number of records shared was excessive; and that the legal basis for the transfer had not been properly established [24]. Information Commissioner Elizabeth Denham stated: "Patients would not have reasonably expected their information to have been used in this way, and the Trust could and should have been far more transparent with patients as to what was happening" [24]. DeepMind acknowledged: "We got that wrong, and we need to do better" [24].

Cambridge researchers Julia Powles and Hal Hodson, writing in "Health and Technology" (2017), called the arrangement "inexcusable" and a "cautionary tale" for the NHS and the healthcare industry [23]. The ethical failure was not that AI was used to detect kidney injury, a clinically valuable application. The failure was that the architecture of the partnership treated patient data as a commercial resource to be extracted rather than a personal asset to be stewarded with explicit consent. The case foreshadowed every subsequent debate about AI data partnerships: who owns health data, who benefits from its use, and what patients have a right to know about how their most sensitive information is deployed.

In October 2021, a legal claim was filed before the High Court of England and Wales seeking compensation for the misuse of private information on behalf of the 1.6 million affected NHS patients [25]. The case remains a defining precedent for AI data partnerships globally.

Case 3: UnitedHealthcare nH Predict (When Accountability Disappears in the Algorithm)

In November 2023, a federal class action lawsuit was filed against UnitedHealth Group and its subsidiary NaviHealth in Minnesota federal court, alleging that the AI algorithm "nH Predict" was used to systematically deny post-acute-care coverage (nursing home stays, home health services) to elderly Medicare Advantage beneficiaries, overriding determinations made by their treating physicians [26]. The complaint alleged that when coverage denials driven by nH Predict were appealed to

federal administrative law judges, approximately 90 percent were overturned, a figure the plaintiffs presented as evidence of the algorithm's fundamental inaccuracy [27].

A U.S. Senate Permanent Subcommittee on Investigations report released in October 2024 added institutional weight to the allegations. The report found that UnitedHealthcare's prior authorization denial rate for post-acute care jumped from 10.9 percent in 2020 to 22.7 percent in 2022 and that the number of denied skilled nursing facility claims in 2022 was nine times higher than in 2019, a period during which the company had "implemented multiple initiatives to automate the process" [6]. CMS responded in February 2024 with an FAQ memo clarifying that AI tools cannot be the sole basis for denial and that algorithms "predicting post-acute care length of stay" cannot be "the sole grounds for terminating services" [28].

In February 2025, a federal judge dismissed five of seven counts but allowed the case to continue breach of contract and breach of implied covenant of good faith and fair dealing claims [29]. As of September 2025, the case remained active, with a judge denying UnitedHealth's bid to limit discovery [29].

These three cases define a taxonomy of ethical failure: algorithmic bias embedded in training data design (Optum); consent and data sovereignty violations in AI partnerships (DeepMind/Royal Free); and accountability erosion in AI-driven clinical and coverage decisions (nH Predict). Each failure mode requires a distinct organizational response, addressed in the "Leadership Implementation" section.

Risks and Limitations

Healthcare AI introduces a set of layered risks that are qualitatively different from the risks of conventional clinical decision support. Understanding them is prerequisite to managing them.

Automation Bias and the Deference Problem

Automation bias, the tendency of humans to over-rely on algorithmic outputs, is not a hypothetical. A November 2024 preprint examining 28 trained pathology experts found a statistically significant 7 percent rate of

automation bias: initially correct evaluations were overturned when erroneous AI advice was introduced [30]. The WHO's 2024 guidance on large multimodal models explicitly warned that large language models (LLMs) "can encourage automation bias by healthcare professionals and patients, whereby errors are overlooked that would otherwise have been identified or difficult choices are improperly delegated" [31]. Two forms are clinically relevant: omission errors, where clinicians fail to notice AI-generated errors; and commission errors, where clinicians defer to AI recommendations despite conflicting clinical evidence [32].

Deskilling: The Google Maps Effect

The theoretical concern about AI deskilling clinical skills became empirical reality in August 2025 with a landmark study published in *The Lancet Gastroenterology and Hepatology*. Across 1,443 non-AI-assisted colonoscopies at four endoscopy centers in Poland, the adenoma detection rate in non-AI-assisted procedures decreased significantly from 28.4 percent before AI exposure to 22.4 percent after, an absolute difference of 6.0 percentage points, among 19 experienced endoscopists over six months [7]. Co-author Romanczyk described the phenomenon as the "Google Maps effect": clinicians who depend on algorithmic guidance can find their independent navigation skills deteriorating [33]. Commentary in the *Lancet* warned that regular AI use may "reduce diagnostic confidence" when AI assistance is unavailable, creating "a new human-AI clinical ecosystem" whose risks must be actively managed [33]. This is the first study to present real-world data on deskilling from clinical AI deployment [7].

The Black Box Problem

Cynthia Rudin's 2019 paper in *Nature Machine Intelligence* made the foundational argument: black box (AI systems whose internal decision-making processes are opaque, meaning clinicians can see what the model recommends but not how or why it reached that recommendation) machine learning models should not be used for high-stakes decisions such as healthcare, because post-hoc explanations of opaque models "is likely to perpetuate bad practices and can potentially cause great harm to society" [34]. Rudin's core empirical finding, that for structured clinical data there is no significant performance difference between complex black box models

and simpler interpretable ones, undermines the usual justification for opacity [34]. The practical consequence of black box AI in clinical care is that physicians can't explain diagnoses or treatment recommendations to patients and can't themselves fully evaluate whether an AI recommendation is appropriate for a specific individual [35]. Without explainability, the AMA's code of medical ethics, built on communication, transparency, and shared decision-making, can't be meaningfully honored. In my experience, when I can't explain a recommendation to a patient in plain language, that recommendation hasn't earned the right to guide their care.

LLM Hallucinations in Clinical Settings

Large language models used in clinical contexts hallucinate: they generate confident, plausible-sounding content that is factually incorrect. A 2025 study in *npj Digital Medicine* evaluating LLM-generated documentation across 450 primary-care consultation transcripts found that 1.47 percent of sentences contained hallucinations, of which 44 percent were classified as "major," meaning they could affect patient diagnosis and management if uncorrected [36]. Major hallucinations appeared most frequently in the "Plan" section of clinical notes, the section containing direct instructions to patients and colleagues [36]. In radiology settings, hallucination rates in experimental systems have been reported at 8–15 percent [37]. Research has established that hallucination "may be an intrinsic, theoretical property of all LLMs," meaning technical mitigation can reduce but not eliminate the risk [36]. The clinical implication is unambiguous: the clinician-in-the-loop is not optional governance theater. It's a patient safety requirement.

Re-Identification and Privacy Risk

The HIPAA Safe Harbor de-identification standard was designed for a world of paper records. A 2019 study published in *Nature Communications* demonstrated that 99.98 percent of patients in a de-identified dataset could be re-identified using only 15 demographic attributes [38]. AI systems that process, analyze, and share derived data (including genomic data, retinal images, and social determinants of health) amplify re-identification risks by enabling triangulation across multiple data sources. University of Michigan law professor W. Nicholson Price II has noted: "AI also makes it easier to

re-identify patients from de-identified data that's been shared by triangulating from different data sources" [39].

The Liability Vacuum

As of early 2026, there are no published judicial opinions that have squarely addressed how the standard of care is altered by the use of an AI system in clinical practice [40]. No settled case law governs whether a physician is liable for following an incorrect AI recommendation, whether a hospital bears responsibility for deploying a biased AI tool, or whether an AI developer can be sued under product liability doctrine for a harm-causing algorithm. This legal uncertainty creates perverse incentives: clinicians may treat AI as confirmatory advice only, neither fully drawing on nor appropriately questioning its outputs, while hospitals lack clear guidance on what constitutes adequate validation, monitoring, and training [40]. Meanwhile, 88 percent of AI vendors impose liability caps on themselves, often limiting damages to monthly subscription fees, while only 17 percent provide warranties for regulatory compliance [41].

HIPAA's Gaps for AI-Specific Protections

HIPAA was enacted in 1996 and has not been updated with AI-specific provisions. As of 2025, HHS has not released any AI-specific HIPAA requirements, though the Office for Civil Rights (OCR, the HHS division responsible for enforcing federal civil rights, privacy, and patient safety laws) has indicated that AI technologies processing protected health information face scrutiny under existing rules [42]. Generative AI tools including standard versions of ChatGPT are not HIPAA compliant because their developers do not enter into business associate agreements with covered entities [43]. The proposed January 2025 update to the HIPAA Security Rule, the first major revision in 20 years, would tighten safeguards for electronic PHI, but it had not been finalized as of early 2026 [42].

Regulatory and Legal Landscape

The regulatory response to healthcare AI bias has been fragmented but accelerating. Federal agencies, state legislatures, and international bodies are all building oversight mechanisms, often without coordination. The

frameworks described in this section represent the current regulatory infrastructure as of early 2026, with particular attention to the mechanisms most relevant to bias, equity, and accountability in clinical AI.

FDA Framework: Scale, Speed, and the 510(k) Question

More than 1,250 AI-enabled medical devices are now authorized for U.S. marketing [1]. In 2024 alone, the FDA authorized a record 168 AI/ML-enabled devices, the most in any single year, with radiology comprising the dominant share [44]. The striking regulatory fact is that 97 percent of these devices were cleared via the 510(k) pathway, which requires only demonstration of substantial equivalence to a prior device, not independent clinical evidence of safety and effectiveness in real-world patient populations [8].

The FDA's framework has evolved materially in recent years. The December 2024 final guidance on Predetermined Change Control Plans (PCCPs), codified under the 2022 Food and Drug Omnibus Reform Act, allows manufacturers to implement pre-authorized AI model modifications without filing new marketing submissions, addressing the mismatch between traditional regulatory review cycles and continuously learning AI systems [45]. The final PCCP guidance notably expanded diversity considerations, requiring manufacturers to account for intended use populations by race, ethnicity, disease severity, gender, and age [45]. In January 2025, the FDA issued draft guidance on AI-Enabled Device Software Functions lifecycle management, acknowledging that modern AI's adaptive capacity complicates oversight designed for static algorithms [46].

In June 2024, the FDA jointly published with Health Canada and the United Kingdom's MHRA guiding principles for transparency of machine learning-enabled medical devices, addressing six dimensions: who receives what information, why, about what, where, when, and how [47]. These are advisory principles, not regulatory requirements. The Trump administration's January 2026 announcement at the Consumer Electronics Show that the FDA would ease oversight of clinical decision support software, with Commissioner Makary calling it "cutting unnecessary regulation," created immediate uncertainty about the durability of transparency expectations for the growing segment of AI tools that meet nondevice criteria [48].

ONC HTI-1: Transparency Inside the EHR

The ONC Health Data, Technology, and Interoperability Final Rule (HTI-1), published in January 2024, represents the most operationally specific transparency requirement currently in force for clinical AI embedded in certified health IT [49]. The rule requires developers who include "predictive decision support interventions" in certified EHR systems to give end users 31 specific "source attributes," covering the intervention's purpose, development details, fairness process, external validation, performance measures, and update schedule [49]. These attributes are organized under the Faves framework: the information must allow users to assess whether the tool is Fair, Appropriate, Valid, Effective, and Safe [49]. HTI-1 also requires Intervention Risk Management (IRM) practices with published summaries covering risk analysis, mitigation, and data governance [49].

HHS Section 1557: Civil Rights Law Reaches Clinical AI

The most legally significant development in healthcare AI equity law is HHS's final rule implementing Section 1557 of the Affordable Care Act, published May 6, 2024, and effective July 5, 2024 [50]. For the first time in U.S. law, nondiscrimination protections were explicitly extended to "patient care decision support tools," including clinical algorithms and AI, used in healthcare settings. Covered entities were required to take "reasonable steps" to identify and mitigate discrimination in their AI and decision support tools, with compliance deadlines for most entities by approximately April 2025 [50]. FDA clearance does not exempt a covered entity from Section 1557: if a cleared AI tool produces discriminatory outcomes, the organization using it remains responsible [50]. Failure to comply can trigger federal enforcement actions and loss of federal healthcare funding.

The Executive Order Pendulum

The Biden Administration's October 2023 Executive Order 14110 contained explicit healthcare AI directives for HHS: establish an AI Task Force, develop a strategy for evaluating AI-enabled health technology quality, and create an AI safety program in partnership with patient safety organizations [51]. On January 20, 2025, President Trump revoked EO 14110 as part of

his first-day executive actions, signing EO 14148 [52]. Trump's January 23, 2025, EO 14179 declared U.S. policy to "sustain and enhance America's global AI dominance" and directed a full review of Biden-era AI policies, taking an explicitly deregulatory stance [52]. In December 2025, a further Trump EO created an AI Litigation Task Force within the Department of Justice with authority to challenge state AI laws deemed inconsistent with federal policy [53].

State Laws: A Fragmented Patchwork

In the absence of full federal legislation, states have moved aggressively. By end of 2025, states had enacted more than 125 laws addressing AI development and use [53]. California Governor Newsom signed SB 1120, the "Physicians Make Decisions Act," on September 28, 2024, prohibiting AI from being the sole basis for denying medically necessary care and requiring final coverage decisions to be made by a physician or licensed clinician competent to evaluate the specific situation [54]. Colorado SB 24-205 (the Colorado AI Act), signed May 2024, requires developers and deployers of "high-risk AI systems," including AI used in healthcare decision-making, to conduct annual impact assessments, implement risk management programs, and give patients the right to appeal adverse AI-based decisions through human review [55].

EU AI Act: The Global Standard-Setter

The EU AI Act entered into force on August 1, 2024, the first full AI regulation in the world [56]. For healthcare, the classification consequence is sweeping: all AI systems used as safety components of products requiring third-party conformity assessment under the Medical Device Regulation (Class IIa and above) are classified as high-risk AI [57]. Since approximately 75 percent of commercial AI medical devices relate to radiology, nearly all classified IIa or higher, the vast majority of medical AI on the European market falls under the high-risk compliance regime [57]. High-risk AI requirements include continuous risk management systems, data governance with bias detection obligations, automatic logging, human oversight mechanisms, accuracy and robustness standards, and mandatory registration in the EU AI Act database [57]. Full compliance is required by

August 2026 for new devices and August 2027 for devices already on the market [56].

When I read this regulatory stack in sequence, I'm reminded of what it felt like to be a physician without a clear framework: you improvise, you make good-faith judgments, and you hope the system catches what you miss. The difference is that when the system misses a patient, that patient suffers the consequence, and these regulations exist precisely because AI-related harms have already occurred at population scale.

Physician Accountability: FSMB's Decisive Ruling

In May 2024, the Federation of State Medical Boards, representing 70 state medical and osteopathic boards, adopted a formal report establishing that physicians bear ultimate responsibility for AI-related clinical errors [5]. The FSMB statement is unambiguous: "Once a physician chooses to use AI, they accept responsibility for responding appropriately to the AI's recommendations." Failure to apply human judgment to AI output is characterized as "a violation of a physician's professional duties" [5]. Physicians are required to be able to provide the rationale for their decisions, whether they followed or deviated from an AI recommendation [5]. This ruling draws a clear line: AI is a tool. The physician is the decision-maker.

Ethical Implications

The ethical questions surrounding AI in healthcare extend beyond bias. They include foundational issues of consent, explainability, liability, and the allocation of moral responsibility when algorithms influence clinical decisions. The sections that follow address the most pressing of these questions, starting with the one closest to the patient encounter.

Informed Consent for AI-Assisted Care

The ethical obligation to obtain informed consent before using AI tools in patient care is now established in professional standards. The FSMB's May 2024 report states explicitly that physicians must receive a patient's consent prior to applying an AI tool to their care and that the duty to explain

diagnosis and treatment options includes the use of AI [5]. The *American Journal of Bioethics* (AJOB) published a 2024 framework identifying three normative functions of a patient's right to notice and explanation for AI use: notification to keep patients informed, increased understanding and trust, and a necessary component of valid informed consent [58]. The AJOB framework argues that if patients would refuse AI involvement in their care when given notice, this is itself sufficient reason why AI involvement falls under the scope of informed consent: withholding notice denies patients the autonomy to decide the nature of their own care [58].

A 2025 scoping review in *BMC Health Services Research* examining 27 studies on digitalized informed consent found that AI-based consent tools are "not yet suitable for use without medical oversight" and are currently best positioned as preparatory or supplementary instruments rather than substitutes for physician conversation [59]. Courts reviewing analogous cases involving new surgical technologies have consistently found physicians liable for consent malpractice when patients were not adequately informed that a technique was new, experimental, or carried unknown risks, a legal precedent with direct implications for AI-assisted clinical encounters [60].

Right to Explanation

The 2025 AMA policy on AI explainability calls for an independent third party, not the developer, to determine whether an algorithm meets the standard of explainability [61]. The AMA's position is clear: intellectual property concerns do not nullify a patient's right to transparency and autonomy in their medical decisions [61]. The General Data Protection Regulation (GDPR, the European Union's comprehensive data privacy law) reinforces this principle at the regulatory level, requiring that algorithmic decision-making be transparent before it is used in patient care [62]. The AJOB framework extends the principle further, proposing that patients have an affirmative right to know not just that AI was used but how it functioned and what limitations it carries [58].

AI Scribes and the Doctor-Patient Relationship

AI ambient documentation tools represent one of the most broadly adopted and patient-visible forms of clinical AI. The TPMG/NEJM Catalyst analysis

of 7,260 Permanente physicians across more than 2.5 million patient encounters found that AI scribes saved an estimated 15,791 hours of documentation time over one year [63]. Physician satisfaction was high: 84 percent reported a positive effect on patient interactions, 47 percent of patients said their doctor spent less time at the computer, and 56 percent reported a positive impact on visit quality [63]. A UCLA randomized trial found that AI scribes delivered measurable reductions in burnout and cognitive workload, though the senior investigator emphasized that "this technology requires active physician oversight, not passive acceptance" [64].

The ethical questions are not resolved by these positive findings. AI scribe tools present consent questions (do patients know their consultation is being transcribed and processed by AI?) and documentation accuracy risks: when AI notes contain errors that physicians fail to catch, those errors enter the permanent medical record and can affect future care. In my experience reviewing AI-assisted clinical notes, the errors that slip through most often occur in the plan section, precisely where Obermeyer's hallucination framework study found major errors most concentrated [36].

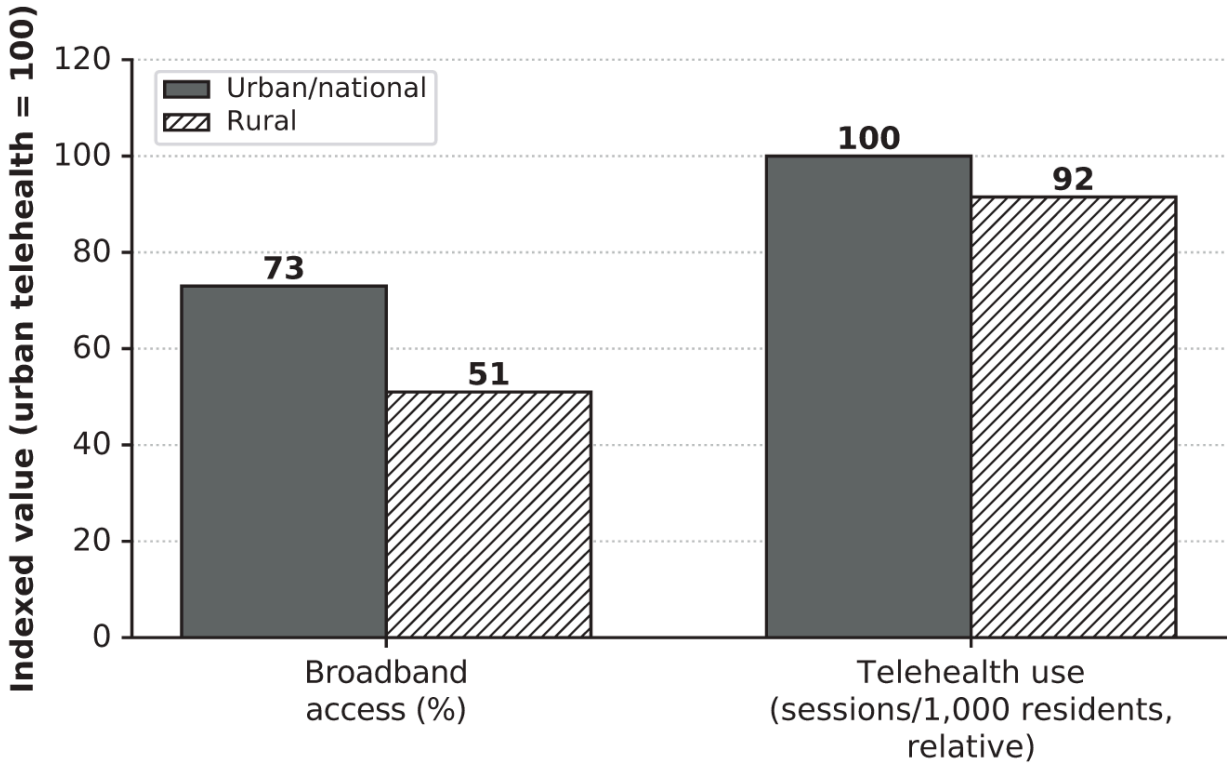
Equity and the Digital Divide

AI-based healthcare tools can create a secondary access inequality that compounds the primary access inequalities they are sometimes intended to address. A national study of nearly 8 million telemedicine sessions found that adults in rural counties used telehealth at significantly lower rates than urban dwellers (approximately 8.5 fewer sessions per 1,000 residents), with broadband access identified as a stronger predictor of telehealth use than almost any other social vulnerability dimension [65]. In identified healthcare desert regions, only 51 percent of households subscribe to fixed broadband compared to 73 percent nationally [65]. Roughly 36 percent of rural students can't afford broadband compared to 28 percent of urban students (see Figure 10.1) [66].

Three pillars must be addressed simultaneously: infrastructure availability, affordability, and digital literacy [67]. Deploying AI health tools without addressing these structural gaps risks creating a two-tier system in which AI accelerates the care quality of well-connected patients while leaving the most vulnerable further behind (see Figure 10.2).

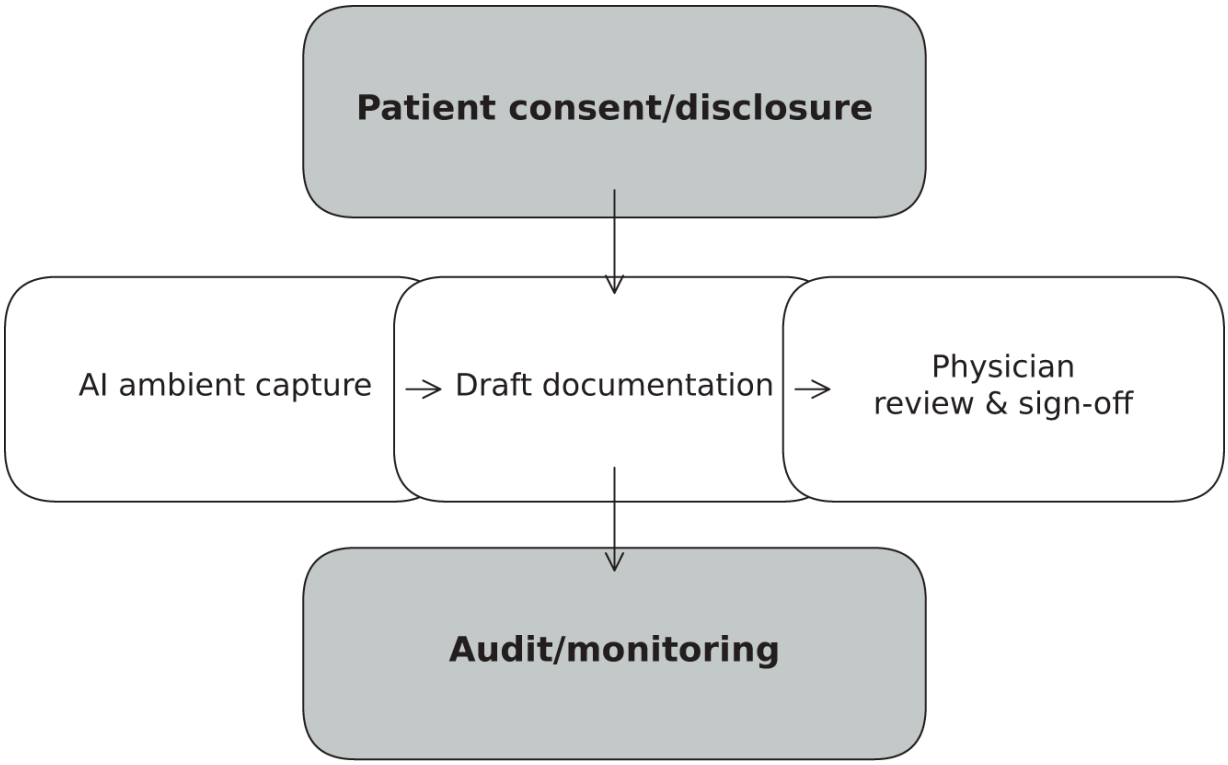
Leadership Implementation Framework

Addressing the ethical challenges described in this chapter requires institutional infrastructure, not just individual awareness. The most effective approach emerging across healthcare organizations is the formation of dedicated AI ethics committees with defined authority, clear membership criteria, and structured review processes.



Source: Manuscript Chapter 10 — 8.5 fewer telehealth sessions per 1,000 rural residents; 51% rural vs 73% national broadband access.

Figure 10.1 The AI adoption equity gap across the United States: rural versus urban broadband affordability and telehealth-use disparities [65,66]



Conceptual diagram synthesised from Chapter 10 (AMA position on consent and audit).

Figure 10.2 Four pillars of ethical AI in healthcare: (1) fairness/bias mitigation, (2) transparency/explainability, (3) accountability/liability, (4) equity of access

Forming an AI Ethics Committee

Every healthcare organization deploying AI at scale requires a standing AI Ethics Committee with genuine authority, not just an advisory role. Effective composition should include:

- Clinical representatives from high-deployment specialties (radiology, pathology, critical care, emergency medicine)
- Informaticists and data scientists with technical AI expertise
- Legal and compliance officers
- Patient and community representatives with explicit attention to demographic diversity
- Ethicists

- A physician executive with direct board accountability

The committee's mandate should cover pre-deployment evaluation of new AI tools, ongoing monitoring of deployed tools for bias and performance drift, incident response, policy development, and annual reporting to the governing board. The Joint Commission/CHAI Responsible Use of AI in Healthcare guidance (September 2025) specifies that governance should create accountability and drive safe AI use, with the governing body receiving periodic updates on AI use, outcomes, and potential adverse events [68]. I've found that ethics committees work only when they have real authority to say no. Advisory panels without veto power are theater (see Figure 10.3).

Bias Detection and Mitigation: A Three-Stage Framework

The literature supports a structured three-stage approach aligned with the AI development lifecycle [69].

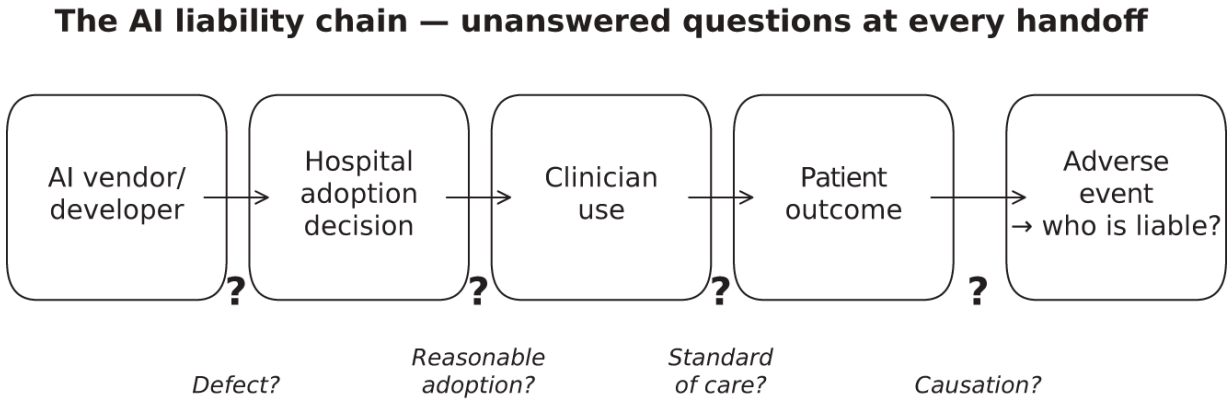


Figure 10.3 The AI liability chain: unanswered accountability questions at every handoff from developer to patient

Preprocessing interventions address training data before model development:

- Auditing demographic representation in training datasets
- Correcting imbalanced data through re-weighting or synthetic augmentation
- Removing biased proxy variables (as in the Optum case, substituting illness severity for costs)

- Documenting data diversity for all deployed tools [69]

In-processing interventions apply during model training:

- Adversarial debiasing techniques that prevent models from inferring sensitive attributes from outputs
- Fairness regularization constraints
- Causal modeling approaches that identify the pathways by which sensitive attributes influence predicted outcomes [69]

Post-processing interventions adjust model outputs after training:

- Group-specific recalibration
- Threshold adjustments
- Equalized odds calibration to ensure equivalent true positive and false positive rates across patient subgroups [69]

Systematic reviews suggest preprocessing approaches currently show the greatest practical effectiveness, while post-processing calibration, if poorly implemented, can worsen prediction errors [69].

Fairness Metrics for Clinical AI

Three foundational fairness metrics should be applied when evaluating clinical AI for bias [70]:

Statistical parity (demographic parity): Requires that model predictions be independent of sensitive attributes, ensuring equal positive outcome rates across patient groups. Best suited for situations where outcome rates should genuinely be equal across groups.

Equal opportunity: Ensures that patients likely to achieve a positive outcome have an equal chance of receiving a positive prediction, regardless of demographic attributes. Appropriate when the goal is equity in identifying true cases.

Equalized odds: Requires both equal true positive rates and equal false positive rates across all patient groups. The most stringent

standard and typically most appropriate for high-stakes diagnostic scenarios [70].

It is mathematically impossible to satisfy all three metrics simultaneously in many real-world settings. The choice of metric is therefore itself an ethical decision, context-dependent, specialty-specific, and requiring explicit deliberation by the AI ethics committee rather than default technical selection [70].

CHAI Blueprint and Joint Commission Standards

The Coalition for Health AI (CHAI), whose 1,240 organizational members include Johns Hopkins University, Mayo Clinic, Stanford Medicine, Google, and Microsoft, released its Assurance Standards Guide v0.3 in June 2024, structuring healthcare AI governance around a six-stage AI lifecycle and five core principles:

- Usefulness/Efficacy
- Fairness and Equity
- Safety and Reliability
- Transparency/Accountability
- and Security/Privacy [71]

CHAI recommends “model cards” as structured disclosure instruments, effectively nutrition labels for AI tools, giving clinicians transparency into model performance, intended use, training data composition, and known limitations [72].

The Joint Commission and CHAI jointly released “The Responsible Use of AI in Healthcare” (RUAIH) in September 2025, the first major governance document from the Joint Commission on AI [68]. Its seven core elements are:

- AI Policies and Governance Structures
- Patient Privacy and Transparency
- Data Security and Data Use Protections

- Ongoing Quality Monitoring
- Voluntary Blinded Reporting of AI Safety-Related Events
- Risk and Bias Assessment
- Education and Training [68]

While currently nonbinding, this guidance is expected to inform future Joint Commission accreditation standards.

Algorithmic Impact Assessments

An algorithmic impact assessment (AIA) is a structured evaluation of an AI system's potential effects, positive and negative, on identified patient populations before deployment. The Ada Lovelace Institute, which designed the first healthcare-specific AIA for the NHS AI Lab's National Medical Imaging Platform (2022), proposed AIAs as a prerequisite for data access, not just post-deployment monitoring [73]. In practice, a healthcare AIA should address:

- The patient populations affected and their demographic composition
- The decision types the AI will influence
- Identified bias risks and proposed mitigation measures
- Consent and transparency obligations
- Monitoring plan and metrics
- Escalation pathways if performance degrades after deployment [73]

The Ada Lovelace Institute notes that AIAs work best in combination with other accountability mechanisms, including audits, transparency registers, and incident reporting, not as stand-alone compliance exercises [73].

Colorado SB 24-205 now requires annual impact assessments for high-risk AI in healthcare [55]. Under HHS Section 1557, covered entities must document "reasonable steps" to identify and mitigate AI discrimination [50]. Organizations should treat AIAs as both an ethical obligation and a legal protection.

Patient Notification and Consent Protocols

Organizations should develop explicit patient notification protocols for AI-assisted care, drawing on FSMB, AMA, and AJOB guidance. At minimum, patients should be informed, in plain language and their preferred language, when AI tools are influencing their diagnosis, treatment recommendations, or coverage determinations. They should be offered the opportunity to request a human review. Documentation of this disclosure should be recorded in the medical record. For AI ambient documentation tools, patients should be informed at the start of the encounter that the consultation is being transcribed and processed by AI, and should have the opportunity to decline. The AMA requires that disclosure occur "in a culturally and linguistically appropriate manner" [74].

Third-Party Auditing and Accreditation

Self-assessment is insufficient for building public trust in clinical AI. URAC has developed a Health Care AI Accreditation program providing third-party review of governance practices, covering safety, fairness, transparency, and accountability [75]. The AMA's 2025 explainability policy calls for independent parties, independent parties such as regulatory agencies or medical societies, to determine whether algorithms meet explainability standards, explicitly excluding developer self-certification [61]. The Stanford HAI FURM framework, operational at Stanford Health Care, gives a model for proactive AI evaluation that assesses utility, ethical implications for patient subgroups, and disclosure obligations before deployment [76]. Organizations should require third-party validation reports as a condition of AI procurement.

I've found in practice that independent audits change the conversation with vendors in ways that internal reviews never do. When an outside party has validated performance claims, the negotiation shifts from credibility to terms, and that shift is where health systems recover real leverage on liability and monitoring obligations.

Vendor Accountability in Contracts

Current market conditions heavily favor vendors: 88 percent impose liability caps limiting damages to monthly subscription fees, and only 17

percent provide warranties for regulatory compliance [41]. Healthcare organizations must rebalance these terms. The Stanford Hai policy brief specifically recommends that licensing agreements require:

- Developers to give information necessary for ongoing risk assessment and monitoring
- Indemnification clauses for errors originating in the model
- Minimum insurance requirements
- Defined responsibilities for post-deployment monitoring [77]

The Joint Commission/CHAI guidance specifies that data use agreements with third-party vendors must define permitted uses, prohibit re-identification of de-identified datasets, require compliance with the organization's privacy standards, and preserve audit rights [72].

Incident Reporting and Response

Healthcare organizations should establish AI-specific incident reporting channels distinct from general adverse event systems. The Joint Commission/CHAI RUAIH guidance promotes voluntary, blinded reporting of AI safety-related events, including unsafe recommendations, major performance degradation, and biased outputs, through existing Patient Safety Organizations (PSOs) and an emerging CHAI registry [68]. ECRI documented in webinar polling that only 7 percent of attendees were aware of AI-related incidents at their facility that had or could have caused patient harm, confirming systematic underreporting [78]. A rapid response protocol for AI-related incidents should include:

- Immediate assessment of patient impact
- AI tool suspension pending investigation if appropriate
- Root-cause analysis distinguishing model failure from implementation failure
- Corrective action documentation
- External reporting to FDA MedWatch when required [78]

Future Outlook

The ethical and regulatory frameworks described in this chapter were largely developed in response to predictive AI, systems that classify, score, and flag. Generative AI introduces a fundamentally different set of risks, and the field is only beginning to grapple with what those risks mean for patient safety, clinical liability, and institutional governance.

Generative AI: New Frontiers, New Risks

Generative AI is creating ethical frontiers that earlier AI ethics frameworks were not designed to navigate. A 2025 systematic review in “Frontiers in Digital Health” identified four primary ethical challenges unique to generative AI in healthcare:

- Systemic bias from non-representative training data
- Unresolved legal liability
- Opacity of complex generative models
- Data privacy risks including the potential for training data memorization and inadvertent patient re-identification [79]

The WHO’s 2024 guidance on large multimodal models, encompassing more than 40 recommendations, signals that the global governance community recognizes these as qualitatively distinct challenges from the narrow AI tools that dominated the FDA device approval pipeline through 2023 [31].

Autonomous AI and the Future of Clinical Decision-Making

The legal and ethical architecture of healthcare AI was constructed around the assumption of a human clinician-in-the-loop who ultimately decides. Autonomous AI systems, those capable of generating and acting on recommendations without human review of each individual output, challenge that architecture at its foundation. The Frontiers in Medicine analysis of liability for AI in medicine distinguishes between assistive AI (where clinicians retain full liability because they can independently

evaluate outputs) and autonomous AI (where creators should assume liability because the human is not positioned to evaluate individual decisions in real time) [80]. As AI systems in radiology, pathology, and ophthalmology move toward autonomous operation, particularly in settings where specialist access is limited, the governance frameworks for consent, accountability, and incident response will require corresponding evolution.

AI in End-of-Life Care

AI mortality prediction systems are being deployed to trigger palliative care discussions, raising ethical questions that go beyond standard bias and accuracy concerns. A 2025 systematic review in the *Interactive Journal of Medical Research* found that AI in palliative care generates algorithmic predictions that can be used without adequate informed consent, introduces bias against marginalized groups when models are trained on non-representative data, and can produce overreliance leading to premature end-of-life planning that fails to account for individual complexity [81]. End-of-life care intersects deeply with cultural preferences around death and dying that Western-dominated training datasets systematically underrepresent [81]. AI ethics committees evaluating end-of-life AI tools must apply heightened scrutiny to patient autonomy, cultural sensitivity, and the distinction between algorithmically predicted trajectories and individual patient values.

Pediatric AI Ethics: The Pearl-AI Framework

Children are systematically underrepresented in healthcare AI development because they represent a smaller, heterogeneous population with fewer chronic diseases, making the commercial return on pediatric AI investment lower [82]. A 2025 systematic review in *npj Digital Medicine* introduced the Pediatrics EthicAl Recommendations List for AI (Pearl-AI) framework, identifying 11 core ethical considerations for AI in child health:

- Nonmaleficence
- Beneficence
- Autonomy
- Privacy

- Justice and transferability
- Transparency and explainability
- Accountability
- Dependability
- Auditability
- Knowledge management
- Trust [82]

The American College of Radiology recommends that all FDA-authorized AI devices include a statement specifying whether the device is authorized for pediatric use [82]. Health systems deploying AI tools should treat the absence of pediatric-specific performance data as a deployment risk, not just a documentation gap.

From Reactive to Proactive Ethics

The dominant model for healthcare AI ethics to date has been reactive: bias is discovered after deployment, consent violations surface through litigation, and accountability gaps become visible only when patients are harmed. The emerging paradigm, embodied in the Stanford FURM framework [76] and the Joint Commission/CHAI governance guidance [68], is proactive: rigorous pre-deployment evaluation, continuous post-deployment monitoring, systematic incident reporting, and regular bias audits conducted throughout the AI lifecycle.

Stanford's FURM framework, operational at Stanford Health Care and supported by the Patient-Centered Outcomes Research Institute, proactively evaluates how new AI tools might widen or narrow healthcare disparities among patient subgroups, recommends safety monitoring processes before deployment, and determines whether and how AI use should be disclosed to patients [76]. Applied to tools ranging from peripheral arterial disease detectors to cancer risk prediction models, FURM represents a mature model of pre-deployment ethics integration that the field needs to widely replicate.

The Legislative Horizon

The 125+ state AI laws enacted by end of 2025 signal legislative intensity without legislative coherence [53]. The case for full federal legislation is strengthening: a single national standard for AI transparency, bias testing, and patient consent would replace a fragmented patchwork that creates compliance complexity for national health systems without necessarily improving patient protection. The Trump December 2025 EO creating a DOJ AI Litigation Task Force to challenge state laws represents one path, federal preemption through executive enforcement, but does not itself create national standards [53]. The alternative, a bipartisan National AI Commission for healthcare, informed by the National Academy of Medicine's May 2025 AI Code of Conduct [83] and the Joint Commission/CHAI RUAIH guidance [68], would produce durable, evidence-based standards aligned with how care is actually delivered. That work has not yet begun in earnest, but the pressure to begin it grows with every new AI deployment and every new documented harm.

Key Takeaways

Bias is structural, not incidental, and must be treated as a governance design requirement. The Optum algorithm used healthcare costs, not race, as its input variable yet produced systematic racial disparity. Bias enters AI through training data choices, proxy variables, and the systemic inequities encoded in healthcare spending patterns, requiring deliberate disaggregated outcome analysis, not assumption-free trust in algorithmic neutrality.

The three failure modes of healthcare AI are distinct and require distinct organizational responses. Algorithmic bias (Optum) requires training data audits and outcome variable scrutiny. Consent and data sovereignty violations (DeepMind/Royal Free) require explicit patient notification protocols and data use governance. Accountability erosion (nH Predict) requires governance structures that maintain identifiable human decision-making responsibility for consequential clinical and coverage decisions.

ECRI's ranking of AI as the top health technology hazard for 2025 signals that governance has not kept pace with deployment. Governance structures are not lagging the technology; they have not yet been built at the scale the technology demands. ECRI's simultaneous ranking of "insufficient governance" as the second-highest patient safety concern makes the organizational priority clear.

Physicians bear ultimate accountability for AI-assisted clinical decisions: there is no shared liability with the algorithm. The FSMB's May 2024 ruling, that physicians are responsible for AI-related errors once they choose to use AI, is now the operative standard for state medical boards. Failure to apply independent clinical judgment to AI outputs is a violation of professional duty.

LLM hallucinations are intrinsic, irreducible, and most dangerous exactly where the stakes are highest. Clinical AI documentation tools hallucinate at rates of 1.47 percent to 15 percent depending on system and setting, with the most dangerous hallucinations appearing in the plan section of clinical notes. The clinician-in-the-loop is a patient safety requirement, not bureaucratic overhead.

FURM, CHAI, and Joint Commission standards give healthcare organizations actionable governance frameworks they don't need to build from scratch. Stanford's FURM, CHAI's Assurance Standards Guide, and the Joint Commission/CHAI Responsible Use of AI in Healthcare guidance provide operational roadmaps for pre-deployment evaluation, ongoing monitoring, incident reporting, and vendor accountability.

Section 1557 and state laws create enforceable legal compliance floors, not aspirational guidelines. HHS Section 1557 (effective July 2024) requires covered entities to take reasonable steps to identify and mitigate AI discrimination, and FDA clearance does not give exemption. California SB 1120 prohibits AI-only clinical denial. Colorado SB 24-205 requires annual impact assessments. Legal compliance requires active bias monitoring, not passive assumption.

The digital divide is a healthcare AI equity problem that compounds existing disparities. Rural patients use telehealth at significantly lower rates than urban patients, and only 51 percent of

households in identified healthcare deserts subscribe to broadband. Deploying AI tools without addressing infrastructure, affordability, and digital literacy gaps risks accelerating care disparities rather than reducing them.

Vendor contracts must be renegotiated to match the governance obligations that law now imposes. With 88 percent of AI vendors capping their own liability at subscription fees while shifting performance and bias risks entirely to healthcare organizations, standard procurement terms are incompatible with current legal and professional standards. Organizations must require performance guarantees, bias testing documentation, audit rights, and indemnification for model errors.

Proactive ethics, not reactive correction, is the emerging standard of care for healthcare AI. The reactive model of discovering bias after deployment and correcting harm after it occurs is no longer adequate given the scale of AI deployment. Pre-deployment algorithmic impact assessments, prospective fairness metric evaluation, continuous post-market monitoring, and systematic incident reporting are the minimum responsible practice.

References

1. Bipartisan Policy Center. FDA oversight: understanding the regulation of health AI tools [Internet]. Washington (DC): Bipartisan Policy Center; 2025 [cited 2026 Jul 4]. Available from: https://bipartisanpolicy.org/issue-brief/fda-oversight-understanding-the-regulation-of-health-ai-tools.

2. Paubox. AMA releases 8-step AI governance toolkit for healthcare providers [Internet]. Paubox; 2025 [cited 2026 Jul 4]. Available from: https://www.paubox.com/blog/ama-releases-8-step-ai-governance-toolkit-for-healthcare-providers.

3. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019; 366(6464):447-53. doi:10.1126/science.aax2342.

4. ECRI. Artificial intelligence tops 2025 health technology hazards list [Internet]. Plymouth Meeting (PA): ECRI; 2024 [cited 2026 Jul 4]. Available from: https://home.ecri.org/blogs/ecri-news/artificial-intelligence-tops-2025-health-technology-hazards-list.

5. Federation of State Medical Boards. Navigating the responsible and ethical incorporation of artificial intelligence into clinical practice [Internet]. Euless (TX): FSMB; 2024 May 2 [cited 2026 Jul 4]. Available from: https://www.fsmb.org/advocacy/news-releases/fsmb-releases-recommendations-on-the-responsible-and-ethical-incorporation-of-ai-into-clinical-practice.

6. Fox Business. UnitedHealthcare accused of relying on AI algorithms to deny Medicare Advantage claims (citing US Senate Permanent Subcommittee on Investigations report) [Internet]. Fox Business; 2024 [cited 2026 Jul 4]. Available from: https://www.foxbusiness.com/markets/unitedhealthcare-accused-relying-ai-algorithms-deny-medicare-advantage-claims.

7. Budzyn K, Romanczyk M, Kitala D, et al. Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study. Lancet Gastroenterol Hepatol. 2025 Oct;10(10):896-903. doi:10.1016/S2468-1253(25)00133-5. PMID: 40816301.

8. Sivakumar R, Lue B, Kundu S. FDA approval of artificial intelligence and machine learning devices for radiology and other specialties: a systematic review. JAMA Netw Open. 2025; 8(11):e2542338. doi:10.1001/jamanetworkopen.2025.42338. PMCID: PMC12595527.

9. American Hospital Association. 4 actions to close hospitals' predictive AI gap [Internet]. Chicago (IL): AHA Center for Health Innovation; 2025 Nov 4 [cited 2026 Jul 4]. Available from: https://www.aha.org/aha-center-health-innovation-market-scan/2025-11-04-4-actions-close-hospitals-predictive-ai-gap.

10. Paubox. AMA releases 8-step AI governance toolkit for healthcare providers [Internet]. Paubox; 2025 [cited 2026 Jul 4]. Available from:

https://www.paubox.com/blog/ama-releases-8-step-ai-governance-toolkit-for-healthcare-providers.

11. Daneshjou R, Vodrahalli K, Novoa RA, et al. Disparities in dermatology AI performance on a diverse, curated clinical image set. Sci Adv. 2022; 8(32):eabq6147. doi:10.1126/sciadv.abq6147.

12. Shabu AM, Ansari MS, Aslam A. A generative AI approach for reducing skin tone bias in dermatological classification [preprint]. arXiv; 2026 Feb. arXiv:2602.14356. Available from: https://arxiv.org/html/2602.14356.

13. Badrie S. Skin tone bias in AI dermatology tools: are we building inclusive systems? [Internet]. RCSIsmj (Royal College of Surgeons in Ireland Student Medical Journal); 2025 [cited 2026 Jul 4]. Available from: https://rcsismj.com/skin-tone-bias-in-ai-dermatology-tools-are-we-building-inclusive-systems-sadie-badrie.

14. Sjoding MW, Dickson RP, Iwashyna TJ, Gay SE, Valley TS. Racial bias in pulse oximetry measurement. N Engl J Med. 2020; 383(25):2477-8. doi:10.1056/NEJMc2029240. PMID: 33326721.

15. Henry NR, Hanson AC, Schulte PJ, et al. Disparities in hypoxemia detection by pulse oximetry across self-identified racial groups and associations with clinical outcomes. Crit Care Med. 2022; 50(2):204-11. doi:10.1097/CCM.0000000000005394. PMID: 35100193.

16. Gottlieb ER, Ziegler J, Morley K, Rush B, Celi LA. Assessment of racial and ethnic differences in oxygen supplementation among patients in the intensive care unit (pulse oximetry bias review). Ann Am Thorac Soc. 2022. Available from: https://www.atsjournals.org.

17. World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance [Internet]. Geneva: WHO; 2021 [cited 2026 Jul 4]. Available from: https://www.who.int/publications/i/item/9789240029200.

18. UNESCO. Recommendation on the ethics of artificial intelligence [Internet]. Paris: UNESCO; 2021 [cited 2026 Jul 4]. Available from:

https://www.unesco.org/en/artificial-intelligence/recommendation-ethics.

19. McCormick TR. Principles of bioethics [Internet]. Seattle (WA): University of Washington Department of Bioethics and Humanities; [cited 2026 Jul 4]. Available from: https://depts.washington.edu/bhdept/ethics-medicine/bioethics-topics/articles/principles-bioethics.

20. Capella S. How does generative AI affect patients' rights? Voices Bioeth. 2025 Sep 28. Columbia University. Available from: https://journals.library.columbia.edu/index.php/bioethics.

21. Healthcare Finance News. Study finds racial bias in Optum algorithm [Internet]. Healthcare Finance News; 2019 [cited 2026 Jul 4]. Available from: https://www.healthcarefinancenews.com/news/study-finds-racial-bias-optum-algorithm.

22. University of Cambridge. DeepMind-Royal Free deal is 'cautionary tale' for healthcare in the algorithmic age [Internet]. Cambridge: University of Cambridge Research; 2017 [cited 2026 Jul 4]. Available from: https://www.cam.ac.uk/research/news/deepmind-royal-free-deal-is-cautionary-tale-for-healthcare-in-the-algorithmic-age.

23. Powles J, Hodson H. Google DeepMind and healthcare in an age of algorithms. Health Technol (Berl). 2017; 7(4):351-67. doi:10.1007/s12553-017-0179-1.

24. Information Commissioner's Office. Royal Free-Google DeepMind trial failed to comply with data protection law [Internet]. Wilmslow (UK): ICO; 2017 Jul [cited 2026 Jul 4]. Available from: https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2017/07/royal-free-google-deepmind-trial-failed-to-comply-with-data-protection-law.

25. Reuters. NHS patients sue Google DeepMind over use of medical records [Internet]. Reuters; 2021 Oct 1 [cited 2026 Jul 4]. Available from: https://www.reuters.com/technology/google-deepmind-facing-lawsuit-over-use-patient-data-2021-10-01.

26. Ross C, Herman B. Patients sue UnitedHealth, alleging AI used to wrongfully deny coverage to elderly [Internet]. STAT News; 2023 Nov 14 [cited 2026 Jul 4]. Available from: https://www.statnews.com/2023/11/14/unitedhealth-navihealth-ai-lawsuit-nursing-home-denials.

27. Ross C, Herman B. Lawsuit update: UnitedHealth AI denial case [Internet]. STAT News; 2024 Feb 29 [cited 2026 Jul 4]. Available from: https://www.statnews.com/2024/02/29/unitedhealth-navihealth-ai-lawsuit-update.

28. Centers for Medicare & Medicaid Services. FAQ on algorithmic decision-making in Medicare Advantage [Internet]. Baltimore (MD): CMS; 2024 [cited 2026 Jul 4]. Available from: https://www.cms.gov/files/document/alhfaq.pdf.

29. Reuters. Federal judge allows UnitedHealth AI lawsuit over Medicare denials to continue [Internet]. Reuters; 2025 Feb 14 [cited 2026 Jul 4]. Available from: https://www.reuters.com/legal/unitedhealth-prevails-some-counts-ai-lawsuit-over-medicare-denials-2025-02-14.

30. Rosbach E, Ganz J, Ammeling J, Riener A, Aubreville M. When two experts disagree: automation bias in pathology with AI-assisted decision-making [preprint]. arXiv; 2024 Nov. arXiv:2411.00998. Available from: https://arxiv.org/abs/2411.00998.

31. World Health Organization. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models [Internet]. Geneva: WHO; 2024 [cited 2026 Jul 4]. Available from: https://www.who.int/publications/i/item/9789240084759.

32. Goddard K, Roudsari A, Wyatt JC. Automation bias: a systematic review of frequency, effect mediators, and mitigators. J Am Med Inform Assoc. 2012; 19(1):121-7. doi:10.1136/amiajnl-2011-000089. PMID: 21685142.

33. BBC News. AI makes doctors worse at spotting cancer, research shows [Internet]. BBC News; 2025 [cited 2026 Jul 4]. Available from: https://www.bbc.com/news/articles/c93jp2g9gwdo.

34. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell. 2019; 1(5):206-15. doi:10.1038/s42256-019-0048-x.

35. Hurley ME, Lang BH, Kostick-Quenet KM, Smith JN, Blumenthal-Barby J. Patient consent and the right to notice and explanation of AI systems used in health care. Am J Bioeth. 2024; 24(10):102-5. doi:10.1080/15265161.2024.2408981.

36. Asgari E, Montana-Brown N, Dubois M, et al. A framework to assess clinical safety and hallucination rates of LLMs for medical text summarisation. NPJ Digit Med. 2025 May 13; 8(1):274. doi:10.1038/s41746-025-01670-7. PMID: 40360677.

37. Templin T, et al. Large language model hallucinations in medical imaging: prevalence and mitigation. Bioengineering (Basel). 2025. PMCID: PMC12729288. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12729288.

38. Rocher L, Hendrickx JM, de Montjoye YA. Estimating the success of re-identifications in incomplete datasets using generative models. Nat Commun. 2019; 10(1):3069. doi:10.1038/s41467-019-10933-3.

39. Price WN 2nd. Artificial intelligence in health care: applications and legal implications. The Scholar: St Mary's Law Rev Race Soc Justice. 2019; 21. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3404059.

40. Duffourc MN, Cahill RA. Legal implications of AI standard-of-care integration on patients and clinicians. NPJ Digit Med. 2025. doi:10.1038/s41746-025-02157-1.

41. Mello MM, Guha N. Understanding liability risk from healthcare AI. Stanford HAI Policy Brief [Internet]. Stanford (CA): Stanford Institute for Human-Centered Artificial Intelligence; 2024 Feb [cited 2026 Jul 4]. Available from: https://hai.stanford.edu/policy/policy-brief-understanding-liability-risk-healthcare-ai.

42. HIPAA Journal. Proposed HIPAA Security Rule update 2025 and AI [Internet]. HIPAA Journal; 2024 [cited 2026 Jul 4]. Available from:

https://www.hipaajournal.com/proposed-hipaa-security-rule-update-2025.

43. HIPAA Journal. Is ChatGPT HIPAA compliant? [Internet]. HIPAA Journal; 2024 [cited 2026 Jul 4]. Available from: https://www.hipaajournal.com/is-chatgpt-hipaa-compliant.

44. Almarie B, Gonzalez-Gonzalez LF, Dos Santos Barbosa LA, Lutz A, Grosse U, Fregni F. Machine learning-enabled medical devices authorized by the US Food and Drug Administration in 2024. Biomedicines. 2025; 13(12). PMCID: PMC12730494. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12730494.

45. US Food and Drug Administration. Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions (final guidance) [Internet]. Silver Spring (MD): FDA; 2024 Dec [cited 2026 Jul 4]. Available from: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial-intelligence.

46. US Food and Drug Administration. Artificial intelligence and machine learning (AI/ML)-enabled medical devices [Internet]. Silver Spring (MD): FDA; 2025 [cited 2026 Jul 4]. Available from: https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices.

47. US Food and Drug Administration, Health Canada, Medicines and Healthcare products Regulatory Agency. Transparency for machine learning-enabled medical devices: guiding principles [Internet]. Silver Spring (MD): FDA; 2024 Jun [cited 2026 Jul 4]. Available from: https://www.fda.gov/medical-devices/software-medical-device-samd/transparency-machine-learning-enabled-medical-devices-guiding-principles.

48. Park A. Trump's FDA chief says he will ease oversight of AI medical devices [Internet]. Time; 2026 [cited 2026 Jul 4]. Available from:

https://time.com/7205924/trump-fda-ces-ai-medical-devices-regulation.

49. Office of the National Coordinator for Health IT / Assistant Secretary for Technology Policy. HTI-1 final rule [Internet]. Washington (DC): ONC/ASTP; 2024 [cited 2026 Jul 4]. Available from: https://healthit.gov/regulations/hti-rules/hti-1-final-rule.

50. Hogan Lovells. HHS extends nondiscrimination requirements to patient care decision support tools in Section 1557 final rule [Internet]. Hogan Lovells; 2024 [cited 2026 Jul 4]. Available from: https://hoganlovells.com/en/publications/hhs-extends-nondiscrimination-requirements-to-patient-care-decision-support-tools-in-section-1557-final-rule.

51. White House. Executive Order 14110 on the safe, secure, and trustworthy development and use of artificial intelligence [Internet]. Washington (DC): The White House; 2023 Oct 30 [cited 2026 Jul 4]. Available from: https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence.

52. National Law Review. Overview of Trump executive orders impacting AI [Internet]. National Law Review; 2025 [cited 2026 Jul 4]. Available from: https://www.natlawreview.com/article/overview-trump-executive-orders-impacting-ai.

53. National Conference of State Legislatures. Artificial intelligence 2025 legislation [Internet]. Denver (CO): NCSL; 2026 [cited 2026 Jul 4]. Available from: https://www.ncsl.org/technology-and-communication/artificial-intelligence-2025-legislation.

54. California Legislative Information. SB 1120: Health care coverage - utilization review [Internet]. Sacramento (CA): California State Legislature; 2024 [cited 2026 Jul 4]. Available from: https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB1120.

55. Colorado General Assembly. SB 24-205: Consumer protections for artificial intelligence [Internet]. Denver (CO): Colorado Legislature; 2024 [cited 2026 Jul 4]. Available from: https://leg.colorado.gov/bills/sb24-205.

56. European Union. The EU Artificial Intelligence Act (full text) [Internet]. artificialintelligenceact.eu; 2024 [cited 2026 Jul 4]. Available from: https://artificialintelligenceact.eu/the-act.

57. van Kolfschooten H, van Oirschot J. The EU Artificial Intelligence Act (2024): implications for healthcare. Health Policy. 2024. PMID: 39244818. Available from: https://pubmed.ncbi.nlm.nih.gov/39244818.

58. Hurley ME, Lang BH, Kostick-Quenet KM, Smith JN, Blumenthal-Barby J. Patient consent and the right to notice and explanation of AI systems used in health care. Am J Bioeth. 2024; 24(10):102-5. doi:10.1080/15265161.2024.2408981.

59. Goldschmitt M, Gleim P, Mandelartz S, Kellmeyer P, Rigotti T. Digitalizing informed consent in healthcare: a scoping review. BMC Health Serv Res. 2025 Jul. PMCID: PMC12225439. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12225439.

60. Duffourc MN, Cahill RA. Legal implications of AI standard-of-care integration on patients and clinicians. NPJ Digit Med. 2025. doi:10.1038/s41746-025-02157-1.

61. American Medical Association. AMA adopts new policy aimed at ensuring transparency in AI tools [Internet]. Chicago (IL): AMA; 2025 Jun 11 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/press-center/ama-press-releases/ama-adopts-new-policy-aimed-ensuring-transparency-ai-tools.

62. Ahmed F, Naz NS, Khan S, Rehman AU, Ismael WM, Khan MA. Explainable artificial intelligence (XAI) in medical imaging (citing GDPR Article 22). BMC Med Imaging. 2026. PMCID: PMC12809972. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12809972.

63. American Medical Association. AI scribes save 15,000 hours and restore the human side of medicine (reporting on NEJM Catalyst study of The Permanente Medical Group) [Internet]. Chicago (IL): AMA; 2025 Jun 12 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/ai-scribes-save-15000-hours-and-restore-human-side-medicine.

64. UCLA Health. UCLA study finds AI scribes may reduce documentation time [Internet]. Los Angeles (CA): UCLA Health; 2025 Nov 26 [cited 2026 Jul 4]. Available from: https://www.uclahealth.org/news/release/ucla-study-finds-ai-scribes-may-reduce-documentation-time.

65. Cummins MR, et al. Social vulnerability, lower broadband internet access, and rurality associated with lower telemedicine use in U.S. counties. JAMIA Open. 2025 Jul 26; 8(4):ooaf056. doi:10.1093/jamiaopen/ooaf056. PMID: 40718759. PMCID: PMC12296349.

66. Graves JM, Abshire DA, Amiri S, Mackelprang JL. Disparities in technology and broadband internet access across rurality: implications for health and education. Fam Community Health. 2021; 44(4):257-65. doi:10.1097/FCH.0000000000000306. PMCID: PMC8373718.

67. Novak S. Bridging rural America's digital divide in health care. J Med Internet Res. 2025 Dec;27:e88833. doi:10.2196/88833. Available from: https://www.jmir.org/2025/1/e88833.

68. American Hospital Association. Joint Commission unveils first AI guidance in partnership with Coalition for Health AI [Internet]. Chicago (IL): AHA; 2025 Sep 18 [cited 2026 Jul 4]. Available from: https://www.aha.org/news/headline/2025-09-18-joint-commission-unveils-first-ai-guidance-partnership-coalition-health-ai.

69. Sasseville M, Ouellet S, Rhéaume C, Sahlia M, Couture V, Després P, et al. Bias mitigation in primary health care artificial intelligence models. J Med Internet Res. 2025; 27:e60269. doi:10.2196/60269. Available from: https://www.jmir.org/2025/1/e60269.

70. Chinta SV, Wang Z, Palikhe A, Zhang X, Kashif A, Smith MA, et al. Fairness in AI healthcare: a survey. PLOS Digit Health. 2025 May. PMCID: PMC12091740. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12091740.

71. Digital Healthcare Law. CHAI releases version 0.3 of its draft guidance for ethical AI in healthcare [Internet]. Digital Healthcare Law; 2024 Jul 11 [cited 2026 Jul 4]. Available from: https://digitalhealthcare.law/2024/07/11/chai-releases-version-0-3-of-its-draft-guidance-for-ethical-ai-in-healthcare.

72. Hooper Lundy. Joint Commission and the Coalition for Health AI (CHAI) release governance frameworks to address AI in health care [Internet]. Hooper Lundy & Bookman; 2025 Sep 25 [cited 2026 Jul 4]. Available from: https://hooperlundy.com/joint-commission-and-the-coalition-for-health-ai-chai-release-governance-frameworks-to-address-ai-in-health-care.

73. Ada Lovelace Institute. Algorithmic impact assessment: a case study in healthcare [Internet]. London: Ada Lovelace Institute; 2022 Feb 8 [cited 2026 Jul 4]. Available from: https://www.adalovelaceinstitute.org/project/algorithmic-impact-assessment-healthcare.

74. American Medical Association. AMA AI principles [Internet]. Chicago (IL): AMA; 2024 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/system/files/ama-ai-principles.pdf.

75. HealthTech Magazines. URAC accreditation addresses the compliance risks of clinical AI [Internet]. HealthTech Magazines; 2026 Jan [cited 2026 Jul 4]. Available from: https://www.healthtechmagazines.com/urac-accrediting-third-party-clinical-ai-for-reduced-liability.

76. Stanford HAI. Making AI work for health care: Stanford's framework evaluates impact [Internet]. Stanford (CA): Stanford Institute for Human-Centered Artificial Intelligence; 2024 Dec 2 [cited 2026 Jul 4].

Available from: https://hai.stanford.edu/news/making-ai-work-for-health-care-stanfords-framework-evaluates-impact.

77. Mello MM, Guha N. Understanding liability risk from healthcare AI. Stanford HAI Policy Brief [Internet]. Stanford (CA): Stanford Institute for Human-Centered Artificial Intelligence; 2024 Feb [cited 2026 Jul 4]. Available from: https://hai.stanford.edu/policy/policy-brief-understanding-liability-risk-healthcare-ai.

78. ECRI. AI resources; ECRI alert on underreporting of AI-enabled device events [Internet]. Plymouth Meeting (PA): ECRI; 2025 [cited 2026 Jul 4]. Available from: https://home.ecri.org/pages/ai-resources.

79. Tung T, Hasnaeen SMN, Zhao X. Ethical and practical challenges of generative AI in healthcare. Front Digit Health. 2025 Nov 17. PMCID: PMC12665710. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12665710.

80. Cestonaro C, Delicati A, Marcante B, Caenazzo L, Tozzo P. Defining medical liability when artificial intelligence is applied in diagnostic imaging. Front Med (Lausanne). 2023 Nov. PMCID: PMC10711067. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC10711067.

81. García Abejas A, Geraldes Santos D, Leite Costa F, Cordero Botejara A, Mota-Filipe H, Salvador Vergés À. Ethical challenges and opportunities of AI in end-of-life palliative care. Interact J Med Res. 2025 May 14. PMCID: PMC12120364. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12120364.

82. Chng SY, Tern MJW, Lee YS, Cheng LTE, Kapur J, Eriksson JG, et al. Ethical considerations in AI for child health and recommendations for pediatric AI (PEARL-AI framework). NPJ Digit Med. 2025 Mar 10. doi:10.1038/s41746-025-01541-1.

83. National Academy of Medicine. AI code of conduct for health and medicine presented in new NAM special publication [Internet]. Washington (DC): NAM; 2025 May 19 [cited 2026 Jul 4]. Available from: https://www.nationalacademies.org/news/ai-code-of-

conduct-for-health-and-medicine-presented-in-new-nam-special-publication.

AI in Emergency Medicine: What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution, First Edition. Harvey Castro.
© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

11
Humans at Heart, AI in Hand

The scar is on my left hand, between the thumb and index finger. It is faded now, but I can trace it with my fingertip and feel the ridge where the tissue healed wrong, healed without help, really, because the doctor who stitched a 12-year-old boy's hand that night didn't bother with anesthesia and didn't bother with care.

I was 12. A rusty, jagged signpost in the neighborhood had cut me open. My mother, an immigrant, a woman who worked three jobs and fed me with food stamps and government cheese, wrapped my hand in a towel and took me to the emergency room. We waited eight hours. When the doctor finally came, he stitched without numbing, spoke without looking, and left without explaining. The pain was real. But what stayed was worse: the absolute certainty that I did not matter. That I was invisible. That the system saw a chart number, not a child.

For years, that experience repelled me from medicine. Why enter a profession that had treated me and my mother like we were disposable? But the scar kept asking a different question. Not why would you go back but what would it take to make sure this never happens to someone else?

When I look at this scar, I say, you know what? It's making sure that we don't do this to other people. That no one else feels invisible like I did. Like my mother did.

That conviction carried me through college on financial aid, through medical school when people who looked like me were rare in the lecture hall, through an emergency medicine residency that pushed my body to its limit, and into the ER, the same kind of place that had failed me, where I swore I would practice differently.

Then the stroke happened. Then the ICU. Then seven days of being a doctor who couldn't speak, trapped in the same system I had dedicated my life to fixing. And when I came out, I knew that dedication alone wasn't enough. Human effort has a ceiling. I had seen it from both sides of the gurney.

AI didn’t enter my life as a fascination. It entered as a necessity. I needed tools that didn’t get tired at 3 a.m., that didn’t miss drug interactions because the chart was buried in a thousand pages of EHR notes, that didn’t let a 12-year-old boy bleed for eight hours because the system was overwhelmed and no one was watching. The calling was never just medicine. It was making medicine work for everyone, the immigrant mother, the latchkey kid, the patient who can’t speak, the doctor who can’t keep up. Human plus AI will beat the best human or the best AI alone. That is the rung I’m reaching for.

Executive Summary

There is a moment most physicians remember: the moment they knew medicine was their calling. For some it arrives in an emergency room, watching a team pull someone back from the edge of death. For others it comes quietly, in the look of a frightened patient who exhales with relief because a trusted physician sat down, made eye contact, and said, I’m going to take care of you. That moment, irreducibly human, impossible to automate, is why this book exists.

The preceding 10 chapters have charted an extraordinary transformation: artificial intelligence detecting cancers invisible to the unaided eye, algorithms predicting sepsis hours before clinical deterioration, ambient scribes restoring evenings and weekends to physicians who had surrendered them to the electronic health record. The evidence is voluminous and, on the most important points, unambiguous. AI works. It finds things we miss. It processes data at scales no human ever could. And it is already saving lives. What the deepest finding of this entire body of research reveals is not just what AI does alone. It’s about what happens when AI and humans work together.

A landmark 2025 study published in the *Proceedings of the National Academy of Sciences* (*PNAS)*, analyzing more than 40,000 diagnoses across 2,133 clinical case vignettes, documented what researchers called *error complementarity*: humans and AI fail in systematically different ways, and when they work as a collective, each compensates for the other’s blind spots [1]. The hybrid human-plus-AI system outperformed both humans alone and AI alone, not marginally but by a meaningful margin. That

finding is the thesis of this final chapter and the thesis of this book. AI's greatest gift to medicine is not the replacement of physicians. It is the return of physicians to the bedside, unburdened from documentation, freed from administrative noise, restored to the work that drew them to medicine in the first place. As Dr. Eric Topol has called it: AI is the gift of time [2].

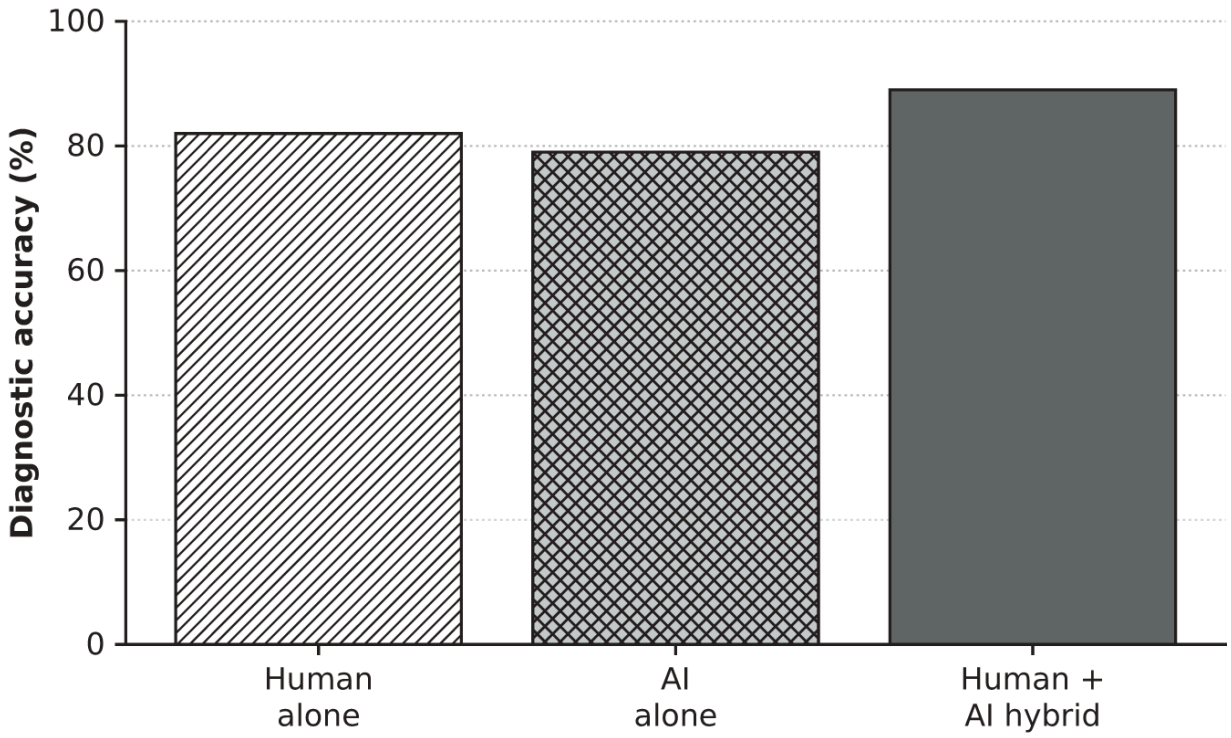
This chapter presents the current evidence, examines institutions that have gotten the balance right, confronts the genuine risks, surveys the regulatory and ethical terrain, and closes with a practical framework for every stakeholder shaping the future of healthcare. The revolution is real. The question before us now is whether we'll direct it wisely, keeping humans at heart, even as we put AI in hand.

Current Evidence

The question of how physicians and AI systems work together has moved from speculation to data. The studies in this section draw from randomized trials, large-scale observational research, and real-world deployment outcomes to establish what human-AI collaboration actually looks like in clinical practice, where it works, and where it breaks down.

The Collaboration Imperative

The science of human-AI collaboration in medicine has matured rapidly (see Figure 11.1). Where early studies asked whether AI could match physicians, the most important recent research asks a more productive question: how do humans and AI together achieve what neither can alone?



Source: PNAS 2025 (40,000+ diagnoses across 2,133 vignettes).
Specific values illustrative; paper reports hybrid system outperforms either alone.

Figure 11.1 Human cognition and AI cognition: the hybrid advantage zone

Evidence from multiple converging lines of inquiry is striking. The 2025 PNAS study led by researchers at the Max Planck Institute for Human Development, the Human Diagnosis Project, and the Italian National Research Council analyzed more than 40,000 diagnoses made by physicians and five leading AI and large language models across 2,133 realistic clinical vignettes [1]. AI collectives alone outperformed 85 percent of individual human diagnosticians on average, yet there remained numerous cases where humans performed better. The mechanism at work was error complementarity: the cases where AI fell short were systematically different from the cases where humans fell short. Hybrid collectives, multiple humans, and multiple AI systems working together produced the most accurate diagnoses in the study. Even adding a single AI model to a group of human diagnosticians noticeably improved diagnostic quality.

I’ve seen this play out firsthand: the cases where I feel most confident are the ones where the data and my clinical read point the same direction. As Senior Research Scientist Stefan Herzog of the Max Planck Institute stated:

“It’s not about replacing humans with machines. Rather, we should view artificial intelligence as a complementary tool that unfolds its full potential in collective decision-making” [1].

This finding has a direct clinical analogue in diagnostic imaging. The Mammography Screening with Artificial Intelligence (MASAI) trial, the first and largest randomized controlled trial of AI in cancer screening, enrolling more than 105,000 women in Sweden, published its final primary results in *The Lancet* in January 2026 [3]. AI-supported screening produced a 29 percent increase in total cancer detection rate, an 80.5 percent sensitivity versus 73.8 percent for standard reading ($p = 0.031$), and a 12 percent reduction in interval cancers detected between screenings. Radiologist workload fell by 44 percent [3]. This is not AI replacing radiologists. It’s AI enabling one radiologist plus AI to match the performance of two radiologists without AI, while detecting more cancers and working more efficiently. The Swedish RCT is the gold-standard demonstration of augmented intelligence in clinical practice.

The “Augmented Intelligence” Consensus

Professional consensus has crystallized around the term the American Medical Association deliberately chose: *augmented intelligence*, not artificial intelligence. The semantic distinction isn’t trivial. The AMA’s policy states that AI tools should “complement and not supplant” physician logic and decision-making [4]. The American College of Physicians echoed this position in its June 2024 policy paper in *Annals of Internal Medicine*, recommending that “AI-enabled technology should be limited to a supportive role in clinical decision-making” [5].

Physician behavior is following the evidence. The AMA’s 2024 Augmented Intelligence Research survey of nearly 1,200 practicing physicians documented that AI use jumped from 38 percent of physicians in 2023 to 66 percent in 2024, a 78 percent increase in a single year [6]. Enthusiasm now exceeds concern for 35 percent of physicians, up from 30 percent the prior year. Three in four physicians believe AI could increase work efficiency, and more than half believe it can reduce stress and burnout [6].

Ambient AI Scribes: Restoring Joy in Medicine

Perhaps the most significant near-term demonstration of augmented intelligence's promise comes not from diagnostic algorithms but from ambient AI documentation and its measurable impact on the physician experience of practice.

The baseline problem is severe. The AMA has documented that for every eight hours of scheduled patient time, office-based physicians spend more than five hours in the electronic health record [7]. More than 45 percent of U.S. physicians report at least one symptom of burnout [8]. Physicians use the phrase *pajama time*, the late-night hours spent finishing notes after the family has gone to bed, as shorthand for an administrative burden that's hollowed out the practice of medicine. Ambient AI scribe data now constitutes some of the most compelling evidence for a technology-enabled solution to this crisis.

At Mass General Brigham, a JAMA Network Open study of 873 clinicians documented a 21.2 percent absolute reduction in burnout prevalence at 84 days of ambient AI scribe use. At Emory Healthcare, a parallel cohort of 557 clinicians experienced a 30.7 percent absolute increase in documentation-related well-being at 60 days [9]. The study's co-senior author, Rebecca Mishuris, MD, chief medical information officer at Mass General Brigham, stated: "There is literally no other intervention in our field that impacts burnout to this extent" [9].

Permanente Medical Group in Northern California published a full 63-week analysis of AI scribe adoption across 7,260 physicians and approximately 2.5 million patient encounters [10]. Physicians using AI scribes saved the equivalent of 1,794 working days in documentation time over one year. Eighty-four percent of surveyed physicians reported the technology had a positive impact on visit interactions, citing reduced mental workload. Forty-seven percent of patients reported their doctor spent less time looking at a computer screen, and two-thirds of patients expressed comfort with the technology [10].

At Cleveland Clinic, more than 4,000 out of 6,000 eligible physicians adopted ambient AI scribes within 15 weeks of availability, a 67 percent adoption rate that reflected organic physician demand rather than administrative mandate. Active users saved 14 minutes per day per clinician

and documented more than one million patient encounters with the system [11].

At the University of Chicago Medicine, a matched-cohort study found ambient AI users spent 8.5 percent less total time in the EHR, with more than a 15 percent drop in time spent composing notes [12].

A six-health-system quality improvement study published in *JAMA Network Open* in October 2025, the first large multicenter evaluation of ambient AI scribes, found burnout rates fell from 51.9 percent to 38.8 percent in 30 days, representing a 74 percent lower odds of experiencing burnout among users [13]. Lee Schwamm, MD, of Yale School of Medicine, a senior author of the study, described the technology as "a really powerful tool to reinvigorate the patient-clinician encounter" [13].

In my experience, nothing gives time back to a physician faster than eliminating the documentation burden. This is augmented intelligence in its clearest form: AI handling the mechanical work of transcription and documentation while the physician turns back to the patient, more present, less exhausted, practicing medicine with restored purpose.

Patient Perspectives: Trust, Consent, and the Human Requirement

Patient perspectives on AI in healthcare are layered in ways that confirm the irreplaceable role of human clinicians. A nationally representative survey of 2,039 U.S. adults published in *JAMIA* found that 30 percent expect AI to improve their access to care and 20 percent expect it to improve their relationship with their doctor, but these expectations were significantly higher among patients who already trust their healthcare providers [14]. Trust in the physician is a prerequisite for trust in AI.

The evidence on AI disclosure is particularly instructive. A June 2025 *JMIR* experimental study found that mentioning a doctor uses AI significantly decreased patient trust, both in the doctor as a person and as a professional ($\beta = -0.40$, $p < 0.001$) [15]. Patients most accept AI when it operates behind the scenes, as a documentation aid, a radiology assistant, or a sepsis monitor, and they least accept it when it becomes the face of their care. A co-design study in the *Journal of Participatory Medicine* quantified this precisely: patients rated AI-as-digital-scribe at 4.24 out of 5 for comfort,

while “virtual human replacing the physician to communicate diagnosis” scored 1.68 out of 5 [16].

What demands attention from every reader is not that patients distrust AI. It’s that they trust AI even when it is wrong. MIT Media Lab researchers presented 300 participants with medical answers from physicians, high-accuracy AI, and low-accuracy AI. AI-generated responses scored higher in trust and satisfaction even when the AI was factually incorrect [17]. This uncalibrated trust, distinct from the informed skepticism that accompanies appropriate use, makes human oversight not just beneficial but essential.

The Empathy Evidence: A Clarifying Lens

A 2023 *JAMA Internal Medicine* study from UC San Diego found that a panel of licensed healthcare professionals, evaluating responses blindly, preferred ChatGPT responses 79 percent of the time over physician responses and rated them 9.8 times more likely to be empathetic [18]. The study authors stated the correct conclusion plainly: “A physician using ChatGPT is the answer for better and empathetic care” [18]. A 2025 meta-analysis in the *British Medical Bulletin* of 15 comparisons confirmed AI’s text-based empathy advantage (SMD 0.87), but the authors noted the critical limitation: all assessments were text-based and could not capture presence, touch, or real-time human connection [19].

The Hojat physician empathy studies, measuring 891 diabetic patients and then 20,961 patients across Italy, documented that physician empathy is significantly associated with better HbA1c and LDL control and with fewer hospitalizations for acute metabolic complications [20]. The quality of the human relationship is itself a clinical intervention. I’ve watched patients relax the moment a physician puts the tablet down and just listens. The goal is the physician equipped by AI, freed from documentation, supported by data, fully present with the patient.

Real-World Case Studies

The programs in this section have moved past proof of concept into sustained clinical operations. They span the largest private health system in the country, the largest government healthcare system, a leading academic

medical center, and resource-limited settings on three continents. Together, they demonstrate what responsible AI deployment looks like when it's done with the patient relationship at the center.

Kaiser Permanente: Building Responsible AI at Scale

Kaiser Permanente's ambient AI documentation program, which we touched on in earlier chapters as a model of enterprise-scale deployment, is the largest deployment of its kind in U.S. healthcare history. The system, powered by Abridge, deployed across all 40 hospitals and more than 600 medical offices in eight states and the District of Columbia, is available to more than 25,000 physicians, advanced practice providers, and pharmacists [21].

The architecture of the program is as instructive as its outcomes. Use is voluntary for clinicians. Patient verbal consent is required before recording begins. Audio recordings are not maintained after note generation. A formal quality assurance framework guided by responsible AI principles governed the rollout, with the explicit design principle articulated by Andrew Bindman, MD, Kaiser Permanente's chief medical officer: "Our clinicians and care teams are the medical decision-makers, not AI" [21].

Published in *NEJM Catalyst*, the outcomes data encompassed 7,260 physicians across 2.5 million encounters over 63 weeks [10]. Documentation time savings were substantial, with the equivalent of 1,794 working days recovered in a year. An independent quality assurance pilot of 63,000 encounters with more than 1,000 physicians rated AI-generated draft notes at an average PDQI score of 4.35 out of 5, with hallucinations documented as "quite rare" [22]. The top-rated domain of the AI-generated notes was "free from bias."

Kaiser Permanente shows that responsible AI deployment, built on consent, transparency, human oversight, and rigorous quality monitoring, can achieve population-scale impact without compromising patient trust.

Department of Veterans Affairs: AI Across 367 Use Cases

The Veterans Affairs healthcare system, the largest integrated health system in the United States, reported 367 AI use cases in its 2025 AI inventory, up from 229 in 2024 and just 40 in 2023 [23]. Of those 367, 107 fall within the

health and medical domain. The breadth of VA AI deployment shows what responsible scaling looks like across a complex institution. It's not a single program. It's a system-wide commitment.

- Storm, the Stratification Tool for Opioid Risk Mitigation, provides one of the most consequential documented outcomes in AI-enabled care. Storm is an AI-assisted clinical decision support system identifying veterans at elevated risk of overdose or suicide. Its implementation is associated with a 22 percent decrease in mortality [23].
- The ReachVet program, running since 2017, has used AI algorithms to identify more than 130,000 veterans at elevated risk, enabling targeted outreach and care before crisis occurs [23].
- VA colonoscopy AI, FDA-approved computer vision devices for polyp detection, has produced a 21 percent increase in the odds of adenoma detection, with higher adenoma detection rates directly associated with reduced late-stage cancer incidence [23].
- VA GPT, the VA's generative AI platform for employees, has onboarded more than 95,000 VA employees. Seventy percent or more of users report improved job satisfaction, saving an average of two to three hours per week [23].

The VA's Chief Technology Officer Charles Worthington framed the organizing principle plainly: thoughtful AI deployment enables "VA's world-class staff to practice at the top of their craft and focus on high-impact work" [23]. These are not AI outcomes. They are human outcomes made possible by AI.

Johns Hopkins Medicine: Closing the Equity Gap

Few real-world case studies more vividly show AI's potential to advance health equity than the Johns Hopkins implementation of autonomous AI for diabetic eye disease screening. Published in *npj Digital Health* in July 2024, the study followed 22,264 patients across more than 30 primary care sites [24]. The results are striking. At AI-enabled sites, adherence to annual diabetic eye testing climbed from 46.4 percent to 64.0 percent, a 17.6 percentage point improvement [24].

The equity data matters even more. Black and African American patients at AI-switched sites saw a 12 percent increase in adherence; the same group at non-AI sites experienced a 0.6 percent decline [24]. The insurance gap between military coverage and Medicaid patients shrank from 32.9 percentage points to 21.1 points. The socioeconomic advantage gap between patients in the top and bottom quartiles of neighborhood disadvantage closed to near zero at AI-switched sites [24].

Diabetic retinopathy is the leading cause of preventable blindness in adults. For decades, screening adherence has been lower in the communities most at risk. A single AI screening tool, deployed within a primary-care workflow, nearly eliminated disparities that generations of public health intervention had failed to close. That's what equitable AI looks like. It's not theoretical.

Zipline, Babyl, and eSanjeevani: AI Where It Is Needed Most

Augmented intelligence is not a phenomenon confined to well-resourced academic medical centers. Three global health programs show AI's capacity to extend the reach of human caregivers to populations that have never had reliable access.

In Rwanda and Ghana, Zipline's AI-powered drone logistics network has reshaped the medical supply chain. A 2022 peer-reviewed study found delivery times shortened by 61 percent versus ground transportation and blood unit expirations reduced by 67 percent [25]. In Ghana, an independent study documented vaccine stock-outs reduced by 60 percent and missed vaccination opportunities cut by 42 percent [25]. As of April 2024, Zipline has completed more than one million commercial deliveries globally, and in February 2026 Rwanda signed a $150 million expansion agreement making it the first country with nationwide autonomous medical logistics coverage [25].

Babyl (formerly Babylon Health, a digital-first primary-care platform using AI-powered triage and virtual consultations originally launched in the United Kingdom) expanded to Rwanda in 2016. By 2021, more than 2.6 million patients were registered, in a country of 12 million with approximately 1,200 doctors, with up to 4,000 consultations per day [26].

Rwanda's Minister of ICT and Innovation articulated the model's logic precisely: "The introduction of Babylon AI allows us to augment human capacity in a chronically understaffed sector" [26].

India's eSanjeevani, the world's largest government-run telemedicine platform, delivered more than 163 million consultations between November 2019 and September 2023 across all 28 states and 8 union territories [27]. In Jharkhand state alone, use grew 1,000-fold, from 50 to more than 50,000 consultations per month. Treatment compliance reached 96.5 percent among patients receiving prescriptions, and 60 percent of patients achieved full recovery following teleconsultation [27].

The thread connecting Kaiser Permanente, the VA, Johns Hopkins, and these global programs is the same: AI deployed with human oversight, equity as an explicit design goal, and the patient relationship preserved at the center of every encounter. These are operational models, and they're showing the rest of the world how it's done.

What strikes me most about these programs is not the technology: it's the governance. Every one of them made human oversight a design requirement, not an afterthought. That's not a constraint on what AI can do; it's the reason these programs work well enough to scale.

Risks and Limitations

The evidence for AI's clinical value doesn't eliminate the risks. In fact, some of the most significant risks emerge precisely because the tools work well enough that clinicians come to depend on them. The sections that follow address the consequences of that dependence, the equity gaps that AI can widen, and the workforce disruptions already underway.

Deskilling: The Most Underappreciated Danger

Every technology that automates a cognitive task carries the risk of eroding the human capacity it replaces. In medicine, that's not an abstract concern. It's clinical.

A 2025 BMJ *Evidence-Based Medicine* editorial warned that overreliance on generative AI risks eroding physicians', particularly medical students', critical thinking skills through automation bias, cognitive offloading, and a

form of dependency the authors called *de-skilling*: the gradual degradation of skills never developed through practice because the algorithm was always there first [28].

The colonoscopy data makes the risk concrete. When AI-assisted adenoma detection is withdrawn from clinical practice, detection rates fall, not just to baseline but below it, as clinicians have come to rely on AI alerts rather than developing independent detection skills. A related finding appears in endoscopy: documented research in Lancet Gastroenterology found adenoma detection rates dropped significantly when AI assistance was removed after a period of regular use, from 28.4 percent to 22.4 percent [28]. Physicians who trained only with AI support had not developed the attentional habits of unaided detection.

Clinicians have started calling this “never-skilling,” the cohort of new trainees who never learn clinical reasoning without algorithmic support, who can’t function at baseline performance when AI is unavailable, and who, over time, may produce a generation of clinicians more dependent on tools than any in history. This is not a reason to reject AI. It is a reason to insist that AI literacy includes knowing when to put the algorithm down. I’ve found that the physicians who stay sharpest are the ones who still work through the differential before checking what the algorithm flagged.

Automation Bias

Automation bias, the tendency to uncritically accept AI outputs without independent reasoning, is among the most consistent findings in human-AI interaction research [29]. A 2024 arXiv study in computational pathology documented a 7 percent automation bias rate: initially correct evaluations overturned by erroneous AI advice [29]. Real-world radiology data showed radiologist override of AI-positive findings declining from 30 percent to 12 percent by year two, a pattern that could reflect appropriate calibration or could reflect declining clinical scrutiny [30].

The distinction matters enormously. Appropriate calibration, learning to trust AI where it is reliable, marks the well-trained AI-era physician. Automation bias, surrendering independent judgment without evaluation, is a patient safety failure. The difference requires training, reflection, and

deliberate cultivation of clinical skepticism as a professional virtue. We don’t get that by accident.

The Trust Paradox

MIT Media Lab findings that patients trust AI medical advice even when the AI is wrong rank among the most alarming data points in the current literature [17]. This uncalibrated trust creates a novel danger: patients who pursue unnecessary care, forgo appropriate treatment, or delay intervention based on erroneous AI output, and who do so with higher confidence than they’d have shown toward erroneous human advice.

AI’s confident communication style, its apparent thoroughness, and the cultural prestige of technology all contribute to this paradox. One practical implication follows clearly: AI systems must only be used alongside medical professionals who can identify, contextualize, and correct AI errors. Deploying direct-to-consumer medical AI without qualified human oversight poses a public safety risk that no efficiency gain justifies.

The Digital Divide and Equity Risk

The same AI-driven clinical tools celebrated in these case studies can widen racial gaps if deployed without equity as an explicit design principle. AHA data documents that large hospitals (over 400 beds) have 90–96 percent AI adoption while small hospitals under 100 beds have only 53–59 percent [31]. Rural hospitals show 56 percent adoption versus 81 percent for urban hospitals; independent facilities show 37 percent predictive AI adoption versus 86 percent for health systems [31]. The communities with the greatest need for AI are the communities least likely to have it deployed. These disparities map directly onto racial demographics: rural and small hospitals disproportionately serve Black, Hispanic, and Native American communities, meaning that the AI adoption gap is, in practice, a racial equity gap.

Training data bias compounds the access gap. The algorithmic bias study we’ve discussed in previous chapters, originally published in *Science* and replicated in *PLOS Digital Health* by Yale researchers, documented that a widely used algorithm prioritized healthier white patients over sicker Black patients for care management because it used cost as a proxy for need, a

proxy that reflects historical underinvestment in minority communities, not current health status [32]. When researchers recalibrated to use direct health indicators, enrollment of high-risk Black patients in care programs nearly tripled, from 17.7 percent to 46.5 percent [32]. Bias in, bias out: AI trained on inequitable historical data will perpetuate and amplify inequity unless designers actively audit and correct for it.

Workforce Disruption

The World Economic Forum's 2025 "Future of Jobs Report" projects that 170 million new jobs will be created in the next five years while 92 million will be displaced, a net gain of 78 million jobs globally [33]. Healthcare is projected to add 3.1 million new nursing positions and 1.6 million new personal care worker positions [33]. The long arc is reassuring. But the transition is painful: 40 percent of employers are considering workforce reductions in areas AI can automate, and the roles most immediately affected, medical coding, clinical documentation, and administrative and transcription functions, are held by real workers whose livelihoods are disrupted now, regardless of aggregate gains by 2030. A just transition requires investment in reskilling for displaced workers, not simply celebration of net job creation figures.

Regulatory and Legal Landscape

The regulatory infrastructure for healthcare AI is being built while the tools are already in clinical use at scale. What follows is the current state of that infrastructure, its most significant gaps, and the specific mechanisms that will shape AI deployment decisions for health systems, physicians, and patients over the next several years.

The Gap Between Innovation and Governance

The United States has reached a moment of genuine regulatory inadequacy. As of January 2026, there are exactly 26 CPT codes for clinical AI solutions, of which only three have achieved Category I status with established reimbursement:

- FFR-CT for cardiovascular risk prediction

- A CT-based cardiovascular risk prediction service
- The autonomous AI tool for diabetic retinopathy detection [34]

Every other FDA-authorized AI solution, and there are more than 1,100 authorized AI/ML-enabled medical devices as of year-end 2024, with 76 percent in radiology [35], navigates a patchwork of Category III codes, carrier-level pricing decisions, and state-level regulatory variations that creates powerful disincentives for responsible deployment.

One principle must anchor any regulatory framework, and the Federation of State Medical Boards has been unambiguous about it: physicians bear ultimate clinical responsibility for decisions made with AI support [36]. This is both ethically correct and, in the current environment, practically burdensome, placing full liability on clinicians for the performance of systems they did not design, cannot fully audit, and often did not select. The AMA's policy recommendation addresses this directly, calling for liability to be appropriately apportioned and for physicians' liability to be limited for AI errors that occur outside their reasonable control [37].

The Federal Standard Imperative

The Healthcare Leadership Council and ZS Associates, publishing a cross-sectoral roadmap in January 2026 developed with experts from 27 organizations, identified governance and regulatory complexity as the primary barrier to responsible AI deployment. Their recommendations are direct:

- Establish centralized federal legislation
- Modernize FDA authority for AI-enabled software as a medical device
- Clarify accountability across all stakeholders
- Ensure liability protections for responsible AI use [38]

Every major professional organization in American medicine, including the AMA, the ACP, and the AHA, has echoed the call for a unified federal AI healthcare standard. The current fragmentation, state-by-state law variations, overlapping federal agency jurisdictions, no unified

reimbursement pathway, does not protect patients. It protects the status quo, which benefits no one.

International Models Worth Studying

The European Union's AI Act, which came into force in 2024, classifies healthcare AI as high-risk and subjects it to conformity assessment, transparency requirements, human oversight mandates, and ongoing post-market monitoring. The UK's Medicines and Healthcare products Regulatory Agency (MHRA) has operated an AI and software as a medical device sandbox since 2022, enabling iterative validation of AI tools in controlled real-world conditions before full market authorization [36].

Neither model is directly applicable to the U.S. healthcare context, but both offer useful architecture. The EU's risk-proportionate classification framework and the UK's sandbox approach to iterative real-world validation are evidence-based regulatory innovations that U.S. policymakers should study carefully and adapt.

CMS and Value-Based Alignment

The Centers for Medicare and Medicaid Services Innovation Center has set a goal that aligns powerfully with AI's capabilities: all Medicare beneficiaries with Parts A and B in an accountable-care relationship by 2030 [39]. As of January 2025, 53.4 percent of traditional Medicare beneficiaries, more than 14.8 million people, were in accountable care relationships, representing the largest annual increase since tracking began [39]. AI's demonstrated capacities in population health management, care gap identification, predictive analytics, and administrative efficiency are precisely the tools that value-based care programs require. Aligning reimbursement policy with AI's demonstrated value, rather than forcing AI tools through fee-for-service coding structures designed for a different era, is the regulatory opportunity before us.

Ethical Implications

The ethical questions in healthcare AI center on a single concern: whether these tools will deepen or diminish the human relationship at the core of

medicine. The frameworks in this section draw from philosophy, narrative medicine, and clinical ethics to articulate what's at stake and what's required to protect it.

The I-Thou Imperative

Martin Buber's distinction between I-Thou and I-It relationships, between genuine human encounter and the transactional use of the other as an instrument, has been applied directly to the physician-patient relationship in medical ethics literature [40]. The I-Thou encounter, characterized by reciprocity, presence, and recognition of the patient as a unique human being, is not just a philosophical ideal. It is a clinical intervention: the Hojat empathy studies documented that the quality of the physician-patient relationship, the "empathic engagement" dimension, is independently associated with diabetic outcomes, chronic pain outcomes, and hospitalization rates [20].

The ethical imperative for AI in medicine follows directly: AI must be designed to create more I-Thou encounters, not fewer. The ambient AI scribe that eliminates "pajama time" and allows the physician to arrive at work rested and present, rather than drained and distracted, serves the I-Thou relationship. The algorithm that frees 10 minutes per appointment from EHR documentation, returning that time to the patient, serves the I-Thou relationship. The AI that interposes itself between physician and patient, generating communications the physician does not review, making decisions the physician does not exercise, threatens to convert medicine's most healing encounters into I-It transactions.

Narrative Medicine in the Age of Algorithms

Rita Charon's narrative medicine, defined as "medicine practiced with the narrative competence to recognize, absorb, interpret, and be moved by the stories of illness," was developed explicitly as a counterbalance to depersonalizing, efficiency-focused healthcare [41]. As AI-generated clinical notes become standard documentation practice, physicians risk shifting from being the authors of their patients' stories to being editors of AI-generated summaries. The phenomenological depth of the patient encounter, what the patient is afraid of, what their illness means in the context of their life, and the subtext beneath the stated chief complaint,

can’t be captured in a structured note. It must be heard, held, and responded to by a human being who is genuinely present.

A 2025 paper in medical humanities argues that medical humanities are “all the more essential” in the AI age, not rendered obsolete by it [42]. The AI era requires not less narrative medicine but more: more deliberate cultivation of the human capacity to be present with illness, more explicit training in the skills of listening and witnessing, more institutional commitment to protecting the time and space for genuine human encounter in clinical practice.

Patient Consent and Transparency: Non-Negotiable

The patient comfort data is unambiguous on one point: transparency and consent are the foundations of AI trust. Patients who were told their data would be used with consent showed 64.7 percent comfort with AI, compared to 44.6 percent overall and 52.6 percent discomfort when data were used without consent [43].

Kaiser Permanente and Cleveland Clinic both require patient verbal consent before ambient AI use, and both cite this consent requirement as integral to their responsible AI frameworks [21]. WHO guidance on large multimodal models is explicit: AI in healthcare must involve patients from the earliest stages of design, patients have the right to understand AI’s function, accuracy, limitations, and data usage, and the deployment of AI without informed consent is not just an ethical violation. It is a threat to the therapeutic relationship that makes clinical care effective [44].

The Equity Imperative

WHO’s organizing ethical principle for healthcare AI, from its six-principle framework on ethics and governance, is that AI must be designed for the widest possible equitable use and access [44]. The National Academy of Medicine (NAM) AI Code of Conduct’s opening commitment is to “Advance Humanity”: to align AI with societal and cultural goals and to support independent evaluation and certification with transparent public reporting [45].

Equity is not a secondary consideration to be addressed after deployment. It must be designed in from the beginning: diverse training data, subgroup

performance monitoring, bias audits before and after deployment, and explicit attention to the communities most at risk of being left behind. The Johns Hopkins diabetic eye screening data shows that equity-centered AI design is achievable, but also that it requires active effort [24].

Even at AI-switched sites that dramatically improved adherence for most demographic groups, some subgroups, Asian patients, non-English speakers, Native American and Alaska Native patients, showed minimal improvement, pointing to deep-rooted access and trust barriers that AI alone cannot resolve [24]. Technology is necessary but not sufficient. Equity requires community engagement, cultural humility, and the human relationships that build trust in communities long underserved by healthcare institutions.

Leadership Implementation Framework

Implementing AI responsibly in healthcare isn't the job of any single role. It requires coordinated action from physicians, health system executives, policymakers, and patients, each operating within their sphere of influence. The following framework addresses each group directly, starting with the clinicians who interact with AI tools at the point of care.

For Physicians: The Engaged Clinician

The time for passive observation has passed. With 66 percent of U.S. physicians now using healthcare AI, up from 38 percent just one year prior, physicians who are not actively engaged with AI tools are already outliers [6]. The question isn't whether to engage but how to engage well.

AI literacy is where this starts. The AMA offers seven CME modules through AMA Ed Hub covering AI basics, methodologies, diagnostic applications, clinical decision support, and ethical and legal considerations, designated for AMA PRA Category 1 Credit and co-developed with the University of Michigan Data-MD team [46]. Stanford Medicine's curriculum covers four learning objectives: how AI works, clinical applications, ethical and legal implications, and critical evaluation [47]. Harvard Medical School has woven AI into all core training and offers a

dedicated AI in Medicine PhD track. Johns Hopkins offers a 10-week AI in Healthcare certificate program for practicing clinicians [47].

Knowing limitations matters as much as knowing capabilities: recognizing automation bias, knowing when to override an algorithm, and maintaining clinical reasoning skills that AI can't replace. The physician of the AI era must evaluate AI recommendations with the same rigor applied to any clinical evidence.

In my experience, the physicians who get the most from AI tools are the ones who treat them the way they'd treat a very smart junior colleague: useful, often right, but not infallible, and never allowed to substitute for your own clinical reasoning. That framing matters. It keeps the physician in the decision seat.

Physicians must also become active participants in AI governance. The AMA's eight-step AI governance toolkit, published May 2025, provides a framework for physician engagement at the health system level: establishing governance structures, inventorying AI tools, defining evaluation processes, and creating feedback channels for AI-related patient safety events [48]. The AMA's Chief Medical Information Officer Dr. Margaret Lozovatsky articulated the obligation clearly: "Physicians must be full partners throughout the AI lifecycle, from design and governance to integration and oversight, to ensure these tools are clinically valid, ethically sound and aligned with the standard of care and the integrity of the patient-physician relationship" [48].

Practically, for every physician:

- Demand transparency from AI vendors
- Insist on reviewing CHAI Applied Model Cards ("AI Nutrition Labels") for every tool used in your practice
- Report AI-related adverse outcomes through patient safety event systems
- Pursue continuing AI education with the same commitment you bring to clinical continuing medical education

For Health System Leaders: Governance Before Deployment

Nine in 10 health system leaders are prioritizing digital and AI products to improve care access, according to a Chartis survey of 150 executives [49]. Yet only 61 percent of hospitals using predictive AI validated those tools on local data prior to deployment, and fewer than half tested for bias before going live [50]. The gap between strategic ambition and responsible execution is where patient harm occurs.

The Joint Commission and the Coalition for Health AI (CHAI) jointly released "The Responsible Use of AI in Healthcare" framework in September 2025, the most thorough governance guidance yet produced for health system leaders [51]. Its seven core elements provide the operational blueprint:

- AI policies and governance structures with cross-functional teams (clinical, IT, compliance, ethics, patient representatives)
- Enhanced data protection
- Ongoing quality monitoring pre- and post-deployment
- Voluntary blinded safety reporting through Patient Safety Organizations
- continuous risk and bias assessment [51]
- Workforce training and education requirements
- Transparency and communication with patients about AI use

American Heart Association's Science Advisory, published in Circulation in November 2025, adds four guiding principles for AI tool selection:

- Strategic alignment with care quality mission
- Ethical evaluation including bias assessment
- Demonstrated measurable clinical benefit
- Financial value commensurate with cost [50]

The practical priorities for health system leaders are clear:

- Stand up an AI governance committee before deploying additional tools, not as a bureaucratic layer but as a genuine multistakeholder forum including patient representatives.
- Require local validation of every AI tool on your patient population before deployment.
- Set continuous post-deployment monitoring with clear thresholds for retraining or retiring underperforming tools.
- Invest in physician and staff AI education as a prerequisite for responsible adoption.
- Measure AI's impact using the Quadruple Aim, patient experience, population health, cost reduction, and provider well-being, not efficiency metrics alone.

For Medical Educators: Training for the Human-AI Team

The AAMC published its seven principles for responsible AI use in medical education, organized at the individual, institutional, and community levels, with the foundational principle that human judgment remains essential and that AI must be integrated while keeping a human-centered approach [52]. The Josiah Macy Jr. Foundation's AI in medical education recommendations, five areas spanning curriculum, assessment, governance, and scholarship, provide the strategic framework for institutional implementation [53].

Effective programs integrate AI across the curriculum rather than siloing it in stand-alone courses. Stanford's approach, four core objectives woven throughout medical school for both MD and PA students, creates a scaffold in which AI literacy develops alongside clinical reasoning, not separate from it [47]. Mount Sinai became the first U.S. medical school to provide all MD and graduate students with access to HIPAA-compliant ChatGPT Edu in May 2025, explicitly framing AI use as a skill to be developed under faculty guidance [47].

Medical educators must train for human-AI teaming, not just AI tool use. The Josiah Macy Foundation's ARiSE Research Network, a Harvard-

Stanford collaboration, is conducting national randomized trials to determine which educational strategies improve doctor-AI teamwork in clinical tasks [53]. Simulation platforms like Weill Cornell's MedSimAI, which uses voice-based AI virtual patients to allow students to practice diagnosis and communication with immediate OSCE-aligned feedback, are training the next generation of physicians to maintain the relational skills that AI can't replicate [54].

Educators must also tackle the never-skilling risk directly. Students who complete residency training having never performed a cognitive task without algorithmic support aren't AI-literate; they're AI-dependent. Curricula must deliberately include training in unaided clinical reasoning, differential diagnosis without algorithmic prompting, physical examination without AI interpretive overlays, as a prerequisite for understanding what AI augments and what it cannot substitute.

For Policymakers: Building the Framework That Protects Everyone

The Healthcare Leadership Council's January 2026 cross-sectoral roadmap, developed with 27 organizations, calls for a clear, forward-looking national standard that harmonizes AI regulations and builds trust across all stakeholder groups, especially patients [38]. That's the policymakers' mandate.

Five urgent actions define the policy agenda.

- A full federal AI healthcare legislation package that establishes a unified regulatory framework, preventing the current state-by-state fragmentation that disadvantages patients and responsible developers equally.
- Modernized FDA authority for AI-enabled software as a medical device, with adaptive authorization pathways that can keep pace with iterative AI development without sacrificing post-market monitoring requirements.
- A reimbursement framework that aligns CMS payment with AI's demonstrated value, expanding CPT coverage beyond the current 26

codes, creating pathways that incentivize responsible deployment and measurable outcomes.

- Federally funded AI equity research addressing the digital divide, mandating bias audits as a condition of Medicare and Medicaid reimbursement for AI-enabled services, and ensuring rural and safety-net institutions have access to AI infrastructure.
- Workforce transition funding: as AI displaces roles in coding, documentation, and administrative functions, investment in reskilling and professional development is not optional.

It is the policy measure that determines whether the transition to AI-augmented healthcare is equitable or merely efficient. AMA's position on the 2025 Federal Government AI Action Plan, calling for vigorous testing, liability apportionment, data quality improvement, and coordinated state-federal regulation, provides the professional consensus that policymakers should engage and act upon [37].

For Patients: Know Your Rights

Patients aren't passive recipients of AI-enabled care. They are participants in a transformation that requires their informed engagement. Every patient interacting with AI-assisted healthcare has the right to know when AI is involved in their care, how their data is used, and how to opt out [44]. Specifically, patients should demand:

- Full informed consent disclosure before AI use in their care
- The right to request human-led care for high-stakes decisions including diagnosis, treatment planning, and end-of-life care
- Transparency about what AI systems their providers are using and what those systems are designed to do
- Data protection assurances, including HIPAA compliance, encryption, and control over personal health information
- Assurance that the AI tools used in their care have been validated for bias across relevant demographic groups, including their own

The trust research is clear: patients who are offered informed consent for AI use, who receive transparent explanations, and who retain the option of human-led care show significantly higher acceptance of AI tools than those who encounter AI without explanation [43]. Transparency isn't a compliance burden. It is the foundation of the therapeutic relationship in the AI era. Patient advocacy organizations should be active participants in AI governance at every level, health system, regulatory, and legislative, ensuring that patient interests shape the technology rather than merely receiving it.

For AI Developers: The Obligations of Power

The National Academy of Medicine's AI Code of Conduct for Health and Medicine, published May 2025, establishes six commitments for every AI developer working in healthcare [45]. They deserve to be quoted in full as the organizing principles for responsible development.

- Advance humanity, align AI with societal and cultural goals for health, support independent evaluation and transparent public reporting.
- Ensure equity, use standardized metrics to assess and report bias in data, AI outputs, and AI use.
- Engage impacted individuals, involve all key stakeholders across the AI lifecycle, including patients, clinicians, and communities.
- Improve workforce well-being, provide reskilling and training programs; ensure AI tools reduce, not increase, clinician burden.
- Monitor performance, use standardized quality and safety metrics to assess AI's impact on health outcomes.
- Innovate and learn, support shared learning, transparency, and a national AI research agenda [45].

NAM co-author Kevin B. Johnson, MD, framed the stakes precisely: "The same thing that happened with electronic health records is happening again with AI, everyone's building tools, but there isn't a shared playbook to make sure they're safe, fair, and actually useful" [45]. The Code of Conduct is that playbook. Developers who build healthcare AI without engaging it are not neutral actors in a technology transition. They are making choices,

about whose data trains the models, whose outcomes are prioritized, whose communities benefit and whose are harmed, whether they acknowledge those choices or not.

The CHAI Applied Model Card, the “AI Nutrition Label” that discloses known risks, training data, bias mitigation approaches, and ongoing maintenance information, should be a non-negotiable deliverable for every AI tool seeking clinical deployment [51]. Physicians deserve to know what is in the tools they are asked to use with their patients. That is not too much to ask. It is the minimum.

Future Outlook

The applications described throughout this chapter, from ambient documentation to diagnostic collaboration to global health delivery, represent the current state of the art. What follows is a look at what’s coming next, the technologies and frameworks that will define the next decade of healthcare AI and determine whether the revolution serves all patients or only some of them.

The Vision Taking Shape

The most significant applications of AI in healthcare are not yet in widespread clinical use. They are in early deployment, in promising trials, and in research pipelines, and the direction they are pointing suggests that the next decade will make the last decade of AI development look cautious.

Digital twin technology—virtual patient replicas built from medical images, EHR data, wearables, and genomic information, continuously updated and able to simulate treatment scenarios—is moving from research to clinical application. A November 2025 review in *Frontiers in Digital Health* documented clinical achievements across specialties:

- Digital twin cardiac models predicting pro-arrhythmic drug risks with 89 percent accuracy, outperforming animal models at 75 percent
- Personalized radiotherapy planning for brain tumors that reduced radiation dose by 16.7 percent while maintaining outcomes

- AI-enhanced glucose management increasing time in target range from 80.2 percent to 92.3 percent during exercise in diabetic patients [55]

Duke University's Center for Computational and Digital Health Innovation has created digital twins that compress time, showing how treatments affect future outcomes without exposing patients to risk [55]. The clinical vision is a shift from reactive to proactive, personalized care: physicians who can model the consequences of treatment decisions before making them, for the specific patient in front of them rather than for a statistical population average.

The FURM framework, Findable, Usable, Responsible, and Meaningful, is emerging as the organizing model for responsible adoption decisions. It asks not just whether an AI tool works in a trial but whether it can reach clinicians who need it, integrate into real-world workflows, be deployed with appropriate governance, and deliver outcomes meaningful to patients. Forward-thinking health systems should apply this lens to every AI tool before and after deployment.

Hospital at Home: AI as the Clinical Infrastructure

The Hospital at Home model, acute care delivered in the patient's home, supported by remote monitoring, telehealth, and AI, is one of the most significant structural changes in American healthcare. As of April 2024, CMS had approved more than 320 hospitals in 37 states to provide acute care at home, with estimates suggesting as much as $265 billion in annual Medicare services could shift to home-based settings [56]. The Mass General Brigham analysis of 5,132 Medicare patients receiving hospital-level care at home documented a 0.5 percent in-hospital mortality rate, with outcomes consistent across Black and Latine patients, dual-eligible, and disabled patients, showing that equity-centered design can achieve equity-distributed outcomes [56].

AI is the clinical infrastructure that makes Hospital at Home scalable. Remote vital signs monitoring powered by AI-enabled wearables detects deterioration before it becomes crisis. Ambient documentation reduces clinician administrative burden in home environments where full hospital support infrastructure is absent. Predictive analytics distinguish true deterioration signals from background noise. Patient engagement platforms

deliver personalized educational content aligned with individual clinical profiles. The physician remains the decision-maker. AI enables a physician's reach to extend to the home, without the patient ever leaving it.

Every time I read data on Hospital at Home outcomes, I think about the patients I saw in the ER who would rather have stayed in their own beds. The technology that makes that possible isn't futuristic: it's already built. What's missing is the leadership will deploy it at scale.

Global Health Equity: The Largest Horizon

The most significant future for healthcare AI is global. One hundred twenty-nine countries have now established national digital health strategies, and more than 1,600 government officials representing more than 100 countries have received specialized training in digital health and AI applications under the WHO's global strategy [57].

India's 2025–26 Union Budget allocated $1 billion for AI-driven digital health services [27]. Rwanda has become the first country with nationwide autonomous medical logistics coverage [25]. Dr. Alain Labrique, WHO Director of Digital Health and Innovation, articulated the mandate: "From AI to telehealth, we have the tools; now we must ensure they reach and benefit everyone" [57].

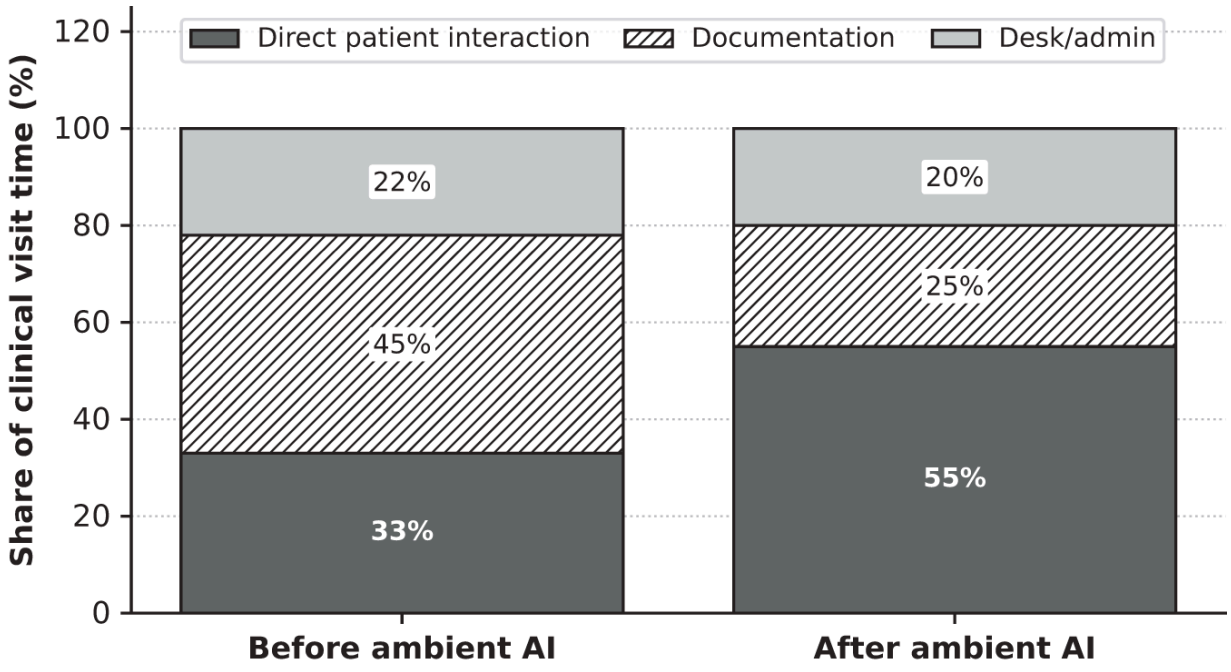
The CAD4TB system achieving 96 percent sensitivity and 85 percent specificity for tuberculosis detection in South Africa and Zambia, the AI chatbot providing pregnancy information in South Africa, the Sierra Leone smartphone app enabling community health workers to detect malaria from blood smear images: these are not proofs of concept [58]. They are the early chapters of a shift in global health equity whose full magnitude will take decades to comprehend.

The Closing Vision

A perspective paper in *Frontiers in Artificial Intelligence* published in February 2026 distilled the challenge precisely: "The question facing medicine is not whether AI can replace clinicians, but whether modern clinical practice can function safely without some form of system-assisted cognition" [59].

Medicine has reached a cognitive inflection point. The exponential growth of biomedical knowledge has outpaced unaided human cognition. AI provides the cognitive scaffolding, speed, pattern recognition, and data synthesis, which allows physicians to navigate this field effectively. But scaffolding is not replacement. The physician brings what AI cannot: accountability, ethical judgment, moral responsibility, and physical presence. The patient in a frightening moment needs both: the pattern recognition that catches what the eye misses and the human presence that makes the unbearable bearable. As Barry Chaiken, MD, has written: “If we do this well, AI will become what it should have been from the start, a partner in healing, not a force that dictates it” (see Figure 11.2) [60].

The revolution is not about technology. It is about what technology makes possible for human care: a medicine more precise, more equitable, more efficient, and more humane than any generation of physicians has ever practiced before (see Figure 11.3).



Conceptual illustration. Directional values consistent with Sinsky et al. (2 hrs doc per 1 hr patient) and ambient-AI deployment outcomes (Chapter 2).

Figure 11.2 The gift of time: from documentation burden to patient presence

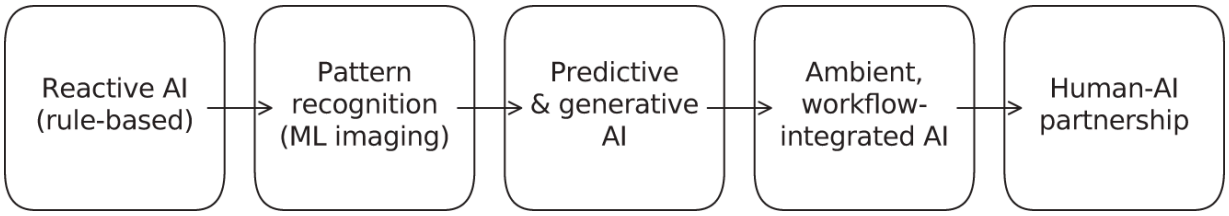


Figure 11.3 The full arc: from reactive AI to human-AI partnership

Key Takeaways

The following are the essential propositions that this revolution demands every stakeholder understand and act upon:

> **Human-AI collaboration outperforms both alone, which is the central clinical finding of this era.** The PNAS 2025 study confirmed error complementarity: humans and AI fail in systematically different ways, and together they achieve what neither can alone. Augmentation, not replacement, is the goal [1].

AI's greatest gift to medicine is time, and the evidence is now documented at scale. Ambient AI scribes are achieving burnout reductions of 13 to 21 percentage points in 30 to 84 days, restoring the time and presence physicians need to practice medicine at its best [9,13].

Patient trust follows human oversight, and that trust is a design requirement, not a communication challenge. Patients most accept AI when it supports rather than replaces human clinicians, when informed consent is offered, and when human-led care remains an option [16]. Organizations that build AI tools around human-centered design earn patient trust; those that deploy AI without transparency risk losing it.

Equity must be designed in from the start; it cannot be corrected after deployment. AI trained on inequitable historical data will perpetuate and amplify inequity. The Johns Hopkins diabetic eye screening program shows that AI can close racial and socioeconomic gaps, but only when equity is an explicit design objective, continuously monitored, and rigorously audited [24].

Deskilling and automation bias are real clinical risks that governance frameworks must actively address. The never-skilling risk is not hypothetical. It demands that medical education preserve unaided clinical reasoning, that governance frameworks include skills maintenance requirements, and that clinical culture reward critical evaluation of AI outputs rather than uncritical acceptance [28].

Governance must precede deployment: the gap between ambition and responsible execution is where patient harm occurs. Only 61 percent of hospitals using predictive AI validated those tools on local data before deployment, and fewer than half tested for bias [50]. The Joint Commission/CHAI responsible use framework, the AMA eight-step governance toolkit, and the NAM AI Code of Conduct collectively define the minimum governance standard. No AI tool should reach patients without meeting it [45,51].

The regulatory framework is inadequate and must change. Twenty-six CPT codes for more than 1,100 authorized AI medical devices is not a reimbursement system. It is an obstacle [34]. A unified

federal AI healthcare legislation framework, aligned with CMS's 2030 accountable-care goals, is the overdue policy priority.

Philosophy of care must guide the technology, not follow it. Martin Buber's I-Thou relationship, Rita Charon's narrative medicine, and the WHO's mandate to advance humanity are not sentimental additions to an otherwise technical project. They are the ethical spine of everything this book has argued [40,41,44]. AI must serve the clinical relationship, not substitute for it.

Global equity is the largest horizon for healthcare AI, and it has barely begun. AI tools detecting tuberculosis in sub-Saharan Africa, extending maternal care to rural India, and delivering vaccines by drone to thousands of health facilities in Ghana are not peripheral stories. They are the most important story this book tells [25–27].

The revolution is a human choice. AI does not have values. It does not have judgment. It does not care. The healthcare revolution described in these pages will be a revolution toward more humane, more equitable, more effective medicine, or it will be something else entirely, depending entirely on whether the humans who design, deploy, govern, and use these tools choose to put patients at the center of every decision. Humans at heart. AI in hand. The rest is ours to do.

Closing Reflection

I want to tell you how I got here. Not the résumé version. The real one.

When I was five years old, my father kidnapped me. He shaved my head so I wouldn't be recognized, enrolled me in a new school under a different name, and disappeared me from my mother's life. She was 21. An immigrant from Colombia. A teenage mother who had given birth to me at sixteen. She had nothing, no money, no legal resources, no power in a system that wasn't built for women like her.

She had determination. For six months, my mother visited every elementary school in New York City. Every single one. She walked in with my photograph, asked if anyone had seen her son, and walked out to the next school. For six months. I cannot fathom the discipline that requires, the

daily decision to keep going when every school says no, when the system offers nothing, when the only resource you have is your own refusal to stop.

She found me. She recruited a friend whose boyfriend drove an 18-wheeler. They came to the apartment where my father was keeping me. My mother grabbed me. My pregnant stepmother grabbed back. There was a tug-of-war, my mother's painted nails drawing blood, my stepmother refusing to let go, a five-year-old boy caught in the middle of two women fighting over his life. My mother won. She threw me into the cab of that semi-truck, and we disappeared. She moved apartments within the hour.

That boy grew up on food stamps. Latchkey kid at four, key on a string around his neck. Government cheese and dried milk. A mother who got her GED while working multiple jobs because stopping was never in her vocabulary. An uncle who showed up with groceries and pretended it was no big deal, so the boy wouldn't know how close to the edge they were.

That boy became a doctor. He walked into the same kind of emergency room that had once treated him like a number and swore it would be different. He made it through residency. He celebrated with his first real shift as an attending. And then his body betrayed him, migraines, a pill, a vasospasm, seven days in the ICU, unable to speak, unable to move, a physician who could hear his own colleagues debating his prognosis and couldn't say a word.

He came back. He always comes back. That's the part of this story that matters.

He came back and started building. Healthcare apps in 54 languages. Coloring book discharge instructions for children. Sign language video guides for patients who are hard of hearing. AI-powered tools that flag drug interactions before they cause vasospasms, that catch cancers before they spread, that detect arrhythmias before they produce strokes. He watched AI do what 17 doctors couldn't: diagnose a child in seconds when years of human effort had failed. He saw a smartwatch save a runner's life. He watched a mammography algorithm catch a cancer that a trained radiologist missed.

And every time he sees it work, every flag, every early detection, every patient who walks out instead of being carried out, he thinks of his mother.

The woman who visited every school in New York City for six months because giving up was not a concept she recognized.

That is the determination healthcare needs to get AI right. Not the determination to be perfect, perfection is a luxury for industries where the stakes are lower. The determination to be relentless. To visit every school. To knock on every door. To build the tool, watch it fail, fix it, and deploy it again. To refuse, absolutely refuse, to let another patient feel invisible.

My mother didn't have AI. She had an 18-wheeler and six months of stubbornness. I have both now, the stubbornness she gave me and the technology this generation is building. The boy in the semi-truck grew up to design systems for the next generation of patients, the next generation of doctors, the next generation of mothers who shouldn't have to save their own lives by accident.

Human plus AI beats the best human or the best AI alone. I learned that from the data. But I believed it long before I saw the evidence. I believed it the night a 21-year-old woman in an 18-wheeler outperformed every system that was supposed to protect her child, and I decided that no system should ever fail that badly again.

References

1. Zöller N, Berger J, Lin I, Fu N, Komarneni J, Barabucci G, et al. Human-AI collectives most accurately diagnose clinical vignettes. Proc Natl Acad Sci U S A. 2025; 122(24):e2426153122. doi:10.1073/pnas.2426153122. Available from: https://www.pnas.org/doi/10.1073/pnas.2426153122 [cited 2026 Jul 4].

2. Topol E. AI in medicine: the gift of time [Internet]. Bethesda (MD): NIH Clinical Center; 2024 [cited 2026 Jul 4]. Available from: https://www.cc.nih.gov/news/2024/nov-dec/eric-topol-ai.

3. ASCO Post Staff. Randomized trial shows AI-supported mammography improves sensitivity and lowers interval cancer rate (MASAI trial) [Internet]. ASCO Post; 2026 Feb 2 [cited 2026 Jul 4]. Available from: https://ascopost.com/news/february-2026/randomized-trial-

shows-ai-supported-mammography-improves-sensitivity-and-lowers-interval-cancer-rate.

4. American Medical Association. AMA issues new principles for AI development, deployment and use [Internet]. Imaging Technology News (ITN); 2023 [cited 2026 Jul 4]. Available from: https://www.itnonline.com/content/ama-issues-new-principles-ai-development-deployment-and-use.

5. American College of Physicians. AI tech should augment physician decision-making, not replace it (Annals of Internal Medicine policy) [Internet]. ecancer; 2024 [cited 2026 Jul 4]. Available from: https://ecancer.org/en/news/24837-ai-tech-should-augment-physician-decision-making-not-replace-it.

6. Henry TA. 2 in 3 physicians are using health AI—up 78% from 2023 [Internet]. Chicago (IL): American Medical Association; 2025 Feb 26 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/2-3-physicians-are-using-health-ai-78-2023.

7. American Medical Association. Digging into the data to cut EHR burdens, drive down burnout [Internet]. Chicago (IL): AMA; 2024 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/digging-data-cut-ehr-burdens-drive-burnout.

8. American Medical Association. National physician burnout survey [Internet]. Chicago (IL): AMA; 2024 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/physician-health/national-physician-burnout-survey.

9. Mass General Brigham. Ambient documentation technologies reduce physician burnout [Internet]. Somerville (MA): Mass General Brigham; 2025 Aug 21 [cited 2026 Jul 4]. Available from: https://www.massgeneralbrigham.org/en/about/newsroom/press-releases/ambient-documentation-technologies-reduce-physician-burnout.

10. Rubio M. Analysis: AI scribes save physicians time, improve patient interactions and work satisfaction (NEJM Catalyst) [Internet]. The Permanente Medical Group; 2025 Jun 9 [cited 2026 Jul 4]. Available from: https://permanente.org/analysis-ai-scribes-save-physicians-time-improve-patient-interactions-and-work-satisfaction.

11. Cleveland Clinic. Less typing, more talking: how ambient AI is reshaping clinical workflow at Cleveland Clinic [Internet]. Consult QD; 2025 Aug 14 [cited 2026 Jul 4]. Available from: https://consultqd.clevelandclinic.org/less-typing-more-talking-how-ambient-ai-is-reshaping-clinical-workflow-at-cleveland-clinic.

12. Niewijk G. Studies suggest ambient AI saves time, reduces burnout and fosters patient connection [Internet]. UChicago Medicine; 2025 Nov 13 [cited 2026 Jul 4]. Available from: https://www.uchicagomedicine.org/forefront/research-and-discoveries-articles/2025/november/ambient-ai-saves-time-reduces-burnout-fosters-patient-connection.

13. Olson KD, Meeker D, Troup M, Barker TD, Nguyen VH, Manders JB, et al. Use of ambient AI scribes to reduce administrative burden and professional burnout. JAMA Netw Open. 2025; 8(10):e2534976. doi:10.1001/jamanetworkopen.2025.34976. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12492056 [cited 2026 Jul 4].

14. Nong P, Ji M. Expectations of healthcare AI and the role of trust: understanding patient views on how AI will impact cost, access, and patient-provider relationships. J Am Med Inform Assoc. 2025; 32(5):795–799. doi:10.1093/jamia/ocaf031. Available from: https://academic.oup.com/jamia/article/32/5/795/8046745 [cited 2026 Jul 4].

15. Chen C, Cui Z. Impact of AI-assisted diagnosis on American patients' trust in and intention to seek help from physicians: randomized survey experiment. J Med Internet Res. 2025; 27:e66083. doi:10.2196/66083.

16. Foresman G, Biro J, Tran A, MacRae K, Kazi S, Schubel L, et al. Patient perspectives on the use of artificial intelligence in clinical care: co-design study. J Particip Med. 2025. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12288699.

17. MIT Media Lab. Patient trust in AI medical advice even when it is wrong [Internet]. Digital Health Insights; 2025 [cited 2026 Jul 4]. Available from: https://dhinsights.org/blog/mit-study-finds-patients-trust-ai-medical-advice-more-than-doctors-even-when-its-wrong.

18. Ayers JW, Poliak A, Dredze M, Leas EC, Zhu Z, Kelley JB, et al. Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. JAMA Intern Med. 2023; 183(6):589–596. doi:10.1001/jamainternmed.2023.1838. Available from: https://jamanetwork.com/journals/jamainternalmedicine/fullarticle/2804309 [cited 2026 Jul 4].

19. Howcroft A, Bennett-Weston A, Khan A, Griffiths J, Gay S, Howick J. Comparing empathy in artificial intelligence and healthcare professionals: a systematic review. Br Med Bull. 2025. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12536877.

20. Hojat M, Louis DZ, Markham FW, Wender R, Rabinowitz C, Gonnella JS. Physicians' empathy and clinical outcomes for diabetic patients. Acad Med. 2011; 86(3):359–364. doi:10.1097/ACM.0b013e3182086fe1. See also: Del Canale S, Louis DZ, Maio V, Wang X, Rossi G, Hojat M, et al. The relationship between physician empathy and disease complications: an empirical study of primary care physicians and their diabetic patients in Parma, Italy. Acad Med. 2012;87(9):1243–1249. doi:10.1097/ACM.0b013e3182628fbf. Available from: https://www.eurekalert.org/news-releases/587799 [cited 2026 Jul 4].

21. Kaiser Permanente. Kaiser Permanente improves member experience with AI-enabled clinical technology [Internet]. Oakland (CA): Kaiser Permanente; 2024 [cited 2026 Jul 4]. Available from: https://about.kaiserpermanente.org/news/press-release-

archive/kaiser-permanente-improves-member-experience-with-ai-enabled-clinical-technology.

22. Kaiser Permanente Division of Research. Quality assurance informs large-scale use of ambient AI clinical documentation (NEJM AI) [Internet]. 2025 [cited 2026 Jul 4]. Available from: https://divisionofresearch.kaiserpermanente.org/quality-assurance-informs-large-scale-use-of-ambient-ai-clinical-documentation.

23. US Department of Veterans Affairs. Building the future: VA's strategy for adopting high-impact artificial intelligence to improve services for veterans [Internet]. Washington (DC): VA; 2025 [cited 2026 Jul 4]. Available from: https://department.va.gov/ai/building-the-future-vas-strategy-for-adopting-high-impact-artificial-intelligence-to-improve-services-for-veterans.

24. Huang JJ, Channa R, Wolf RM, Dong Y, Liang M, Wang J, et al. Autonomous artificial intelligence for diabetic eye disease increases access and health equity in underserved populations. NPJ Digit Med. 2024; 7:196. doi:10.1038/s41746-024-01197-3.

25. Zipline (drone delivery company) [Internet]. Wikipedia; [cited 2026 Jul 4]. Available from: https://en.wikipedia.org/wiki/Zipline_(drone_delivery_company).

26. Babylon launches AI in Rwanda in next step towards digitising healthcare in Rwanda [Internet]. Business Wire; 2021 Dec 3 [cited 2026 Jul 4]. Available from: https://www.businesswire.com/news/home/20211203005293/en/Babylon-Launches-AI-in-Rwanda-in-Next-Step-Towards-Digitising-Healthcare-in-Rwanda.

27. Sood S, Lal K, Bhatia M, Kapoor G, Singh S, Kaushish RK, et al. Adoption and utilization of India's eSanjeevani national telemedicine service. Oxf Open Digit Health. 2025; 3:oqaf025. doi:10.1093/oodh/oqaf025. Available from:

https://pmc.ncbi.nlm.nih.gov/articles/PMC12558045 [cited 2026 Jul 4].

28. BMJ Group. Overreliance on AI risks eroding new and future doctors' critical thinking while reinforcing existing bias [Internet]. BMJ Group; 2025 [cited 2026 Jul 4]. Available from: https://bmjgroup.com/overreliance-on-ai-risks-eroding-new-and-future-doctors-critical-thinking-while-reinforcing-existing-bias.

29. Rosbach E, Ganz J, Ammeling J, Riener A, Aubreville M. When two experts disagree: automation bias in pathology with AI-assisted decision-making [preprint]. arXiv. 2024. arXiv:2411.00998. Available from: https://arxiv.org/abs/2411.00998.

30. Goldsmith-Pinkham P, Tan C, Zentefis AK. Financial arbitrage among artificial intelligence agents [preprint]. arXiv. 2026. arXiv:2601.13379. Available from: https://arxiv.org/abs/2601.13379.

31. American Hospital Association. 4 actions to close hospitals' predictive AI gap [Internet]. AHA Center for Health Innovation Market Scan; 2025 Nov 4 [cited 2026 Jul 4]. Available from: https://www.aha.org/aha-center-health-innovation-market-scan/2025-11-04-4-actions-close-hospitals-predictive-ai-gap.

32. Cross JL, Choma MA, Onofrey JA. Bias in medical AI: implications for clinical decision-making. PLOS Digit Health. 2024; 3(11):e0000651. doi:10.1371/journal.pdig.0000651. Available from: https://journals.plos.org/digitalhealth/article?id=10.1371/journal.pdig.0000651 [cited 2026 Jul 4].

33. World Economic Forum. The Future of Jobs Report 2025 [Internet]. Geneva: WEF; 2025 [cited 2026 Jul 4]. Available from: https://www.weforum.org/publications/the-future-of-jobs-report-2025/. Secondary summary: 5 things HR needs to know from WEF's 2025 Future of Jobs Report. UNLEASH; 2025. Available from: https://www.unleash.ai/artificial-intelligence/5-things-hr-needs-to-know-from-wefs-2025-future-of-jobs-report.

34. Loud G. Paying for AI in U.S. health care [Internet]. Washington (DC): Bipartisan Policy Center; 2026 Mar 16 [cited 2026 Jul 4]. Available from: https://bipartisanpolicy.org/issue-brief/paying-for-ai-in-u-s-health-care.

35. FDA AI/ML medical devices growth analysis [Internet]. MedTech Dive; 2024–2025 [cited 2026 Jul 4]. Available from: https://www.medtechdive.com/news/fda-ai-medical-devices-growth/728975.

36. Pham T. AI in healthcare: physician responsibility and international regulatory models. Biomedicines. 2024. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12076066.

37. American Medical Association. AMA position on 2025 federal government AI action plan [Internet]. Chicago (IL): AMA; 2025 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/practice-management/digital-health/ama-position-2025-federal-government-ai-action-plan.

38. Healthcare Leadership Council. Healthcare sector calls for unified federal framework to unlock AI's full potential for patients [Internet]. HLC/ZS; 2026 [cited 2026 Jul 4]. Available from: https://www.hlc.org/news/healthcare-sector-calls-for-unified-federal-framework-to-unlock-ais-full-potential-for-patients.

39. Centers for Medicare & Medicaid Services. CMS Innovation Center strategy refresh: strategic direction white paper [Internet]. Baltimore (MD): CMS; 2025 [cited 2026 Jul 4]. Available from: https://www.cms.gov/priorities/innovation/strategic-direction-whitepaper.

40. Cohn F. Existential medicine: Martin Buber and physician-patient relationships. J Contin Educ Health Prof. 2001; 21(3):170–181. doi:10.1002/chp.1340210308. Available from: https://pubmed.ncbi.nlm.nih.gov/11563223 [cited 2026 Jul 4].

41. Schwartz D, Lanphier E. Medicine, narrative, and artificial intelligence [Internet]. German National Library (Deutsche Nationalbibliothek);

2025 [cited 2026 Jul 4]. Available from: https://d-nb.info/1373778113/34.

42. Paglialonga M, Simonetti C. Medical humanities in the age of artificial intelligence. Monash Bioeth Rev. 2025. PMID: 40553218. Available from: https://pubmed.ncbi.nlm.nih.gov/40553218.

43. Chandrasekaran R, Takale L, Moustakas E. Listening to patients' voices on the use of AI in health care: cross-sectional study. J Med Internet Res. 2025; 27:e77501. doi:10.2196/77501. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12680129 [cited 2026 Jul 4].

44. World Health Organization. WHO releases AI ethics and governance guidance for large multi-modal models [Internet]. Geneva: WHO; 2024 Jan 18 [cited 2026 Jul 4]. Available from: https://www.who.int/news/item/18-01-2024-who-releases-ai-ethics-and-governance-guidance-for-large-multi-modal-models.

45. National Academy of Medicine. Health care artificial intelligence code of conduct for health and medicine [Internet]. Washington (DC): NAM; 2025 [cited 2026 Jul 4]. Available from: https://nam.edu/our-work/programs/leadership-consortium/health-care-artificial-intelligence-code-of-conduct.

46. American Medical Association. Doctors use AI dramatically—here's the CME they need [Internet]. Chicago (IL): AMA Ed Hub; 2024 [cited 2026 Jul 4]. Available from: https://www.ama-assn.org/membership/continuing-medical-education-cme-online/doctors-use-ai-dramatically-here-s-cme-they-need.

47. Three medical schools pioneering an AI-augmented future (Stanford, Harvard, Mount Sinai) [Internet]. ETC Journal; 2026 Feb 8 [cited 2026 Jul 4]. Available from: https://etcjournal.com/2026/02/08/three-medical-schools-pioneering-an-ai-augmented-future.

48. American Medical Association. 8 steps to position your health system for AI success (AMA STEPS Forward) [Internet]. Chicago (IL): AMA; 2025 [cited 2026 Jul 4]. Available from: https://www.ama-

assn.org/practice-management/digital-health/8-steps-position-your-health-system-ai-success.

49. Chartis. Health system leaders prioritize AI amid healthcare challenges (survey of 150 executives) [Internet]. Healthcare Dive; 2025 [cited 2026 Jul 4]. Available from: https://www.healthcaredive.com/news/health-system-leaders-prioritize-ai-healthcare-challenges-chartis/804539.

50. Jain SS, Goto S, Hall JL, Khan SS, MacRae CA, Ofori C, et al. American Heart Association science advisory on artificial intelligence governance in cardiovascular care. Circulation. 2025. PMID: 41208719. Available from: https://pubmed.ncbi.nlm.nih.gov/41208719.

51. Joint Commission and Coalition for Health AI (CHAI) release governance frameworks to address AI in health care [Internet]. Hooper Lundy & Bookman; 2025 [cited 2026 Jul 4]. Available from: https://hooperlundy.com/joint-commission-and-the-coalition-for-health-ai-chai-release-governance-frameworks-to-address-ai-in-health-care.

52. Association of American Medical Colleges. Principles for the responsible use of artificial intelligence in medical education [Internet]. Washington (DC): AAMC; 2025 [cited 2026 Jul 4]. Available from: https://www.aamc.org/about-us/mission-areas/medical-education/principles-ai-use.

53. Josiah Macy Jr. Foundation. AI in medical education recommendations [Internet]. New York (NY): Macy Foundation; 2025 [cited 2026 Jul 4]. Available from: https://macyfoundation.org/our-grantees/ai-in-med-ed.

54. Cornell Bowers CIS. Medical students use AI to practice communication skills (MedSimAI) [Internet]. Ithaca (NY): Cornell University; 2025 [cited 2026 Jul 4]. Available from: https://bowers.cornell.edu/news-stories/medical-students-use-ai-practice-communication-skills.

55. Khoshfekr Rudsari H, Tseng B, Zhu H, Song L, Gu C, Roy A, et al. Digital twins in healthcare: a comprehensive review. Front Digit Health.

2025. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12671388.

56. Mass General Brigham. Study of national data demonstrates the value of acute hospital care at home (Levine DM, et al. Ann Intern Med) [Internet]. Somerville (MA): Mass General Brigham; 2024 Jan 8 [cited 2026 Jul 4]. Available from: https://www.massgeneralbrigham.org/en/about/newsroom/press-releases/study-of-national-data-demonstrates-the-value-of-acute-hospital-care-at-home.

57. World Health Assembly extends global digital health strategy through 2027 [Internet]. HLTH; 2025 May 26 [cited 2026 Jul 4]. Available from: https://hlth.com/insights/news/world-health-assembly-extends-global-digital-health-strategy-through-2027-2025-05-26.

58. Maruta T. Applications of artificial intelligence in African healthcare diagnostics (CAD4TB, malaria, cervical cancer): a review. JSM Regen Med Tissue Eng. 2024. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC12505451.

59. AI will not replace doctors but augment clinical cognition: the augmented physician (Front Artif Intell) [Internet]. Devdiscourse; 2026 [cited 2026 Jul 4]. Available from: https://www.devdiscourse.com/article/technology/3814246-ai-will-not-replace-doctors-but-augment-clinical-cognition.

60. Chaiken BP. Human oversight and conscience in healthcare AI [Internet]. barrychaiken.com; 2025 Oct [cited 2026 Jul 4]. Available from: https://barrychaiken.com/archives/2025/10/human-oversight-conscience-healthcare-ai.

AI in Emergency Medicine: What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution, First Edition. Harvey Castro.
© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Index

Numerics

21st Century Cures Act, 92, 147, 180

510(k) clearance process, FDA, 4, 10, 66, 92, 112, 120–121, 138, 147, 148, 186, 206, 275, 284

A

AAM. *See* Advance Alert Monitor, Kaiser Permanente

AAMC. *See* Association of American Medical Colleges

Abbott, 140

Abridge, 34–35, 209, 221, 242, 315–316

ACCESS model, CMS, 215, 248, 260

Acclarent, 63

accountability, 217–218, 238, 252, 273–274, 280, 286, 290, 321

accountable-care relationships, 219, 322

ACR. *See* American College of Radiology

ACR Recognized Center for Healthcare-AI (Arch-AI), 121, 122

Acute Hospital Care at Home initiative, CMS, 220

Ada Lovelace Institute, 293

adenoma detection rate, during colonoscopy, 86, 109, 116, 119, 281, 316, 319

ADHD. *See* attention-deficit/hyperactivity disorder

Advance Alert Monitor (AAM), Kaiser Permanente, 209–210, 219, 242

Advocate Health, 220

Affordable Care Act, section 1557 of, 285, 293

agentic AI, 238, 260
AIA. *See* algorithmic impact assessment
AI Action Plan (2025), 72, 329
AI Code of Conduct (2025), NAM, 298, 324, 329–330
AI divide, 216–217, 238, 259
Aidoc, 98
AI-ECG screening, Mayo Clinic, 91, 96
AI-Enabled Device Software Functions, FDA, 72, 284
AI ethics committees, 290, 296
AI Litigation Task Force, 285, 298
AI/ML-Based Software as a Medical Device Action Plan, FDA, 15
AI Risk Management Framework, NIST, 72, 253
AIR-SP (AI Research Screening Platform), 210
AI Steering Committee, 254
AJOB. *See* *American Journal of Bioethics*
alert fatigue, 9, 13, 39, 62, 64, 146, 245
algorithmic accountability, 217–218
algorithmic bias, 12, 17, 54–55, 150, 203, 212–213, 217, 273, 320
  and cardiovascular AI, 59–60
  DeepDerm, 57, 276
  Optum, 56, 245, 252, 274–276, 278–279
algorithmic fairness, 68
algorithmic impact assessment (AIA), 293
algorithmic justice, 68
algorithmic limitations, of AI mental health tools, 179–180
algorithm vigilance/algorithmovigilance, 96, 122

All of Us Research Program, NIH, 72, 99

AlphaFold, 21–22, 221

Alphv. *See* Blackcat (ransomware group)

AMA. *See* American Medical Association

AMA Ed Hub, 325

Ambience Healthcare, 44, 243

ambient AI documentation, 6–7, 10–11, 22, 29, 206, 209, 218, 219, 221, 313–314, 323

  ethical implications, 41–43, 288, 293

  future outlook, 45–47

  JAMA Network Open study, 34–35, 38

  Kaiser Permanente, 10–11, 22, 29, 34–37, 42, 44, 209, 241–242, 258, 287–288, 315–316

  leadership implementation framework, 43–45

  physician time allocation before and after, 29

  RCTs, 33–34

  regulatory and legal landscape, 40–41

  risks and limitations, 37–39

ambient clinical intelligence, 221

American Board of Surgery, 123

American College of Physicians, 218, 312–313

American College of Radiology (ACR), 93, 97, 121, 122, 297

American Heart Association, 154, 327

American Journal of Bioethics (AJOB), 287, 293

American Medical Association (AMA), 4, 14, 31, 35, 38, 68, 70, 93, 121, 150, 207, 216, 237, 238, 251–252, 275, 278, 282, 287, 293, 294, 312, 313, 321, 325

American Nurses Foundation, 31

American Psychological Association (APA), 183, 185, 186

American Society of Nephrology, 69

Americans with Disabilities Act, 150

AMIE, Google, 45, 222

anxiety disorders, 170, 171, 176, 190

APA. *See* American Psychological Association

App Evaluation Model, APA, 186

Apple Heart Study, 137, 139–140

Apple Watch, 85, 136, 138, 141–147, 154, 155

Arch-AI. *See* ACR Recognized Center for Healthcare-AI

ARiSE Research Network, 328

Artificial Intelligence Act (AI Act, 2024), EU, 13, 15–16, 22, 66, 71, 93, 122, 214–215, 249, 286, 322

Ascension Health, 212

Assess-AI registry, ACR, 93, 97, 122

Association of American Medical Colleges (AAMC), 31, 205–206, 218, 220, 327

Assurance Standards Guide, CHAI, 70, 292

ASSURE trial, 46

atrial fibrillation, 8, 84–85, 136–140, 145, 150, 155

attention-deficit/hyperactivity disorder (ADHD), 12, 184

augmented intelligence, 312–314, 317–318

automation bias, 55, 64, 90, 95, 118, 281, 319

autonomous AI systems, 296, 317

autonomous medical logistics, 317

autonomous surgical robots, 112, 124

B

Babyl, 318

BCG. *See* Boston Consulting Group

beneficence (ethics), 277–278

Bertelsmann Foundation, 211

BetterHelp, 178–179, 184, 187

bias, AI, 252. *See also* algorithmic bias; automation bias

amplification of, 17–18

and cardiovascular AI, 59–60

de-biasing, 69–70

detection/mitigation of, 290–291

and diversity of AI development teams, 71

evaluation of, 238

Standing Together consensus recommendations, 70–71

structural nature of, 60

Biden, Joe, 66, 72

Bindman, Andrew, 242, 316

bioethics, 94, 277–278

BiomedGPT, 98, 124

BiomedParse, 114, 124

BJC Health System, 243–244

Black Basta (ransomware group), 212

black box problem, 13, 90, 282

Blackcat (ransomware group), 211, 245

blood pressure monitoring, 143, 154

blood smear photos, analysis of, 117

Blueprint for Trustworthy AI, CHAI, 70

Boston Consulting Group (BCG), 208, 239, 260

bounded justice, 68

brain CT, 114

Brainomix 360, 5

breast cancer screening, 5, 46, 83, 87, 89–91, 108, 109, 113, 117–119, 126, 210, 312

breast tomosynthesis, 89–90

Brigham and Women's Hospital, 64

broadband affordability, 288, 289

build vs buy complexity, 246–247

burnout. *See* physician burnout

C

CAD4TB, 117, 332

California AB 3030, 67, 182, 217

California SB 1120, 249, 285

calorie estimation, 146

cancer

  detection of, 46, 83–84, 86, 91, 99, 108, 109, 113, 117–119, 126, 210, 276

  overdiagnosis of, 89

  skin, 56–58, 276

  Watson for Oncology, 8–9, 61

capital requirement for enterprise AI systems, 213–214

cardiac MRI analysis, 8

cardiovascular AI, 59–60
diagnostic tools, 92
FFR-CT cardiac assessment AI, 210
prediction of cardiovascular disease, 84–85
Care platform, Aidoc, 98
CB Insights, 208
CE marking, European Union, 148–149
Center for Computational and Digital Health Innovation, Duke University, 331
Center for Health AI, 243–244
Centers for Medicare & Medicaid Services (CMS), 92, 93, 173, 181, 205, 207, 215, 219, 220, 249, 260, 280, 322, 328, 331
CEO. *See* chief executive officer
Cerca. *See* Coalition to End Racism in Clinical Algorithms
Cerebral, 178–179, 184, 187
CGMs. *See* continuous glucose monitors
CHAI. *See* Coalition for Health AI
Chandra, Rohit, 242, 243
Change Healthcare, 211–212, 245–246
change management, 153, 246
Character.AI, 177–178, 182, 183
Character Technologies, 178
Charon, Rita, 323
chatbots, 169, 171–173, 175–179, 182, 183, 185–188
ChatGPT, 178, 283, 315
ChatGPT Edu, 327
ChatGPT Health, 180

Chestify AI, 117
chest X-ray AI, 113, 120–121
chief AI officer, 259–260
chief executive officer (CEO), 19–20, 238–239, 241, 252–253
chief information officer (CIO), 20, 253
chief medical officer (CMO), 21, 241, 253–254
CHIME, 253
cholecystectomy, 112, 124
Cigna, 67
CIO. *See* chief information officer
Cleveland Clinic, 44–45, 96, 146, 209, 242, 313–314, 324
CLIA. *See* Clinical Laboratory Improvement Amendments
clinical accuracy of wearables, 144–145
clinical AI deployment
  ethical implications, 69–70
  failures, real-world cases, 61–63
  future outlook, 71–73
  leadership implementation framework, 70–71
  regulatory and legal landscape, 65–68
  risks and limitations of, 63–65
clinical judgment, erosion of, 5, 14, 90
Clinical Laboratory Improvement Amendments (CLIA), 92, 98
clinical leadership, 241
clinical notes, hallucinations in, 37–38, 282
clinical research enterprise, 209
clinical validation of AI tools, 21, 65

clinician adoption, 278. *See also* liability
of AI, 4–5, 14, 207, 213, 237, 238
of wearables, 153
cloud-based AI tools, 20
CMO. *See* chief medical officer
CMR Surgical, 111
CMS. *See* Centers for Medicare & Medicaid Services
Coalition for Health AI (CHAI), 16, 22, 41, 44, 70, 97, 218, 290, 292, 294, 297, 326, 330
Coalition to End Racism in Clinical Algorithms (Cerca), 69
COLO-DETECT trial, 116
colonoscopy, 86, 110, 116, 119, 281, 316, 319
Colorado SB 24-205, 249, 285, 293
Colorado SB 169, 67
colorectal cancer, 116
colorectal surgery, 112, 123
CoMET, 222
commission errors, 281
CommonSpirit Health, 219
Commonwealth Fund, 31
Communications Decency Act, section 230 of, 182–183
community engagement, 71
complete heart block, 142
COMPOSER system, 7
computed tomography angiography (CTA), 115
computer vision algorithms, 210

consent. *See* informed consent
continuous glucose monitors (CGMs), 137, 140–141, 154
convolutional neural networks, 84, 115
cost-effectiveness
  of AI diagnostics, 91
  of clinical AI interventions, 214
  of robotic surgery, 123
COVID-19 pandemic, 30, 58–59
CPT. *See* current procedural terminology codes
credentialing, for robotic surgery, 123
Crisis Text Line, 179, 183
C-Suite AI adoption, 236–237, 326–327
  AI investment, 240
  CEO strategic frameworks, 238–239
  clinical leadership, 241
  and ethical implications, 250–255
  and FDA authorization, 240–241
  future outlook, 259–260
  leadership implementation framework, 252–259
  real-world case studies, 241–244
  and regulatory and legal landscape, 247–250
  and risks and limitations of healthcare AI, 244–247
  state of, 237–238
CTA. *See* computed tomography angiography
cuffless blood pressure monitors, 154
current procedural terminology (CPT) codes, 93, 114, 151, 215, 248, 321

cyberchondria, 146
cybersecurity, 20, 120, 203, 211–212, 245–246

D

DALL·E, 6
Dartmouth College, 169, 171, 189
data-based operation and communication center (DOCC), SMC, 211
data breaches, 211–212, 245–246
data colonialism, 68
data commodification, 178–179, 184, 279–280
data governance, 186–187, 217
Data Protection Act (1998), United Kingdom, 279
dataset shift, 65
Datasheets for Datasets, 71
Davidoff, Ramin, 37
da Vinci 5 (Model Is5000), 111, 119, 125
DAX Copilot, 10–11, 22, 33, 34, 38, 42, 43, 45, 219, 221
DaylightRx, 173
Ddi dataset, 69
de-biasing, 69–70
DeepDerm, 57, 276
deep learning models, 13, 90, 118
  ECG-based, 208
  race prediction by, 60
deference problem, 14, 64, 90, 109, 118, 281
Delft Imaging, 117
Deloitte, 207–208, 236, 238, 239, 260

demographic bias in AI imaging, 120
denials, insurance coverage, 67, 280
Department of Veterans Affairs (VA), 45, 316–317
depression, 169–172, 176, 184, 189, 190
dermatology
  AI diagnostic tools in, 8
  skin cancer diagnosis, 56–58, 276
deskilling, AI-driven, 5, 14, 38, 109, 119, 281, 318–319
DETECT study, 116
developers, AI, 329–330
development team diversity, 71
Dexcom, 140
diabetes, 140–141
diabetic retinopathy screening, 8, 65, 88, 217, 317, 324
diagnostic errors, 80, 82, 95, 222
diagnostic partnership, 45–46
diagnostic reasoning, and LLMs, 86
digital divide, 14, 18, 149–150, 183–184, 288, 320
Digital Health Advisory Committee, FDA, 181
Digital Health Index (2024), 211
digital phenotyping, 189
digital therapeutics (DTx), 173–174, 181
digital twins, 221–222, 331
disclosure, 69, 121, 122, 182, 293, 314
DOCC. *See* data-based operation and communication center, SMC

documentation burden, 6–7, 28–29, 32–33, 333. *See also* ambient AI documentation

Dragon Copilot, 6, 7, 22, 40

drift detection, 72, 97–98

drug discovery, 21–22, 221

DTx. *See* digital therapeutics

Duke University, 331

Duke University Hospital, 85, 96

duty to detect, 93–94

Dyania Health, 209

E

eating disorders, 171–172, 175–176, 190

ECG

  app, Apple Watch, 138, 140, 145, 147

  interpretation, 8, 84

  patches, 141

ECRI, 3, 18, 42, 65, 252, 275, 278, 295

EEOC. *See* Equ al Employment Opportunity Commission

efficiency gains, physician productivity, 42, 43

eGFR race correction, 55, 69–70

EHDS. *See* European Health Data Space

Ehrenfeld, Jesse, 251

EHRs. *See* electronic health records

electronic health records (EHRs), 6–7, 9, 28–29, 35–36, 211, 238, 284. *See also* ambient AI documentation

adoption of, 30

data quality, 60

Epic, 62

gaps, 13

integration, 20, 61, 152, 189

time spent by physicians on, 32

vital signs, 60

Emory Healthcare, 258, 314

empathy, 315, 323

employers

mental health platforms sponsored by, 176

wearables mandated by, 150

EndeavorRx, 173

end-of-life care, 296

energy expenditure estimation, 146

enterprise AI, 202–203, 315

cost barriers, 212–214

real-world cases, 208–210

Epic Sepsis Model (ESM), 9, 13, 21, 62, 90, 245

Epic Signal, 34

Epic Systems, 222

Equal Employment Opportunity Commission (EEOC), 150

equalized odds (fairness metric), 292

equal opportunity (fairness metric), 291–292

EQUAL trial, 137, 140, 155
equity, 18, 212–213, 324–325. *See also* ethics
access to AI diagnostics, 94–95
access to surgical robots, 119
AI adoption gap, 39, 42–43, 288, 289, 320
AI divide, 216–217, 238, 259
and AI mental health tools, 183–184
diabetic eye testing, 317
Ethics-Evidence-Equity framework, AMA, 68
and false positives, 89
global health, 332
Hospital at Home model, 220, 331–332
research, 328
section 1557 of Affordable Care Act, 285, 293
and wearables, 150, 152, 155
error complementarity, 310, 312
eSanjeevani, 318
ESM. *See* Epic Sepsis Model
Essence Healthcare, 144
Estate of Lokken v. UnitedHealth Group, 67
Estonia, national health system of, 211

ethics, 273–275, 286–288. *See also* algorithmic bias
AI-assisted surgery and imaging, 121–122
AI mental health tools, 183–185
ambient AI documentation, 41–43, 288, 293
autonomous AI systems, 296
clinical AI deployment, 69–70
failures, case studies, 278–281
frameworks, 277–278
generative AI, 295
healthcare AI, 17–19
human-AI collaboration, 323–325
and leadership implementation framework, 288, 290–295
and pace of physician AI adoption, 278
pediatric AI, 296–297
pillars of ethical AI, 289
predictive diagnostics, 93–95
proactive, 297
reactive, 297
scale of deployment of AI tools, 275
wearables, 149–150
"Ethics and Governance of Artificial Intelligence for Health," WHO, 250, 277
Ethics-Evidence-Equity framework, AMA, 68
EU. *See* European Union
European Health Data Space (EHDS), 217

European Union (EU)
2030 eHealth targets, 211
AI Act (2024), 13, 15–16, 22, 66, 71, 93, 122, 214–215, 249, 286, 322
GDPR, 122, 179, 287
Medical Device Regulation (EU 2017/745), 148–149, 286
Everhart, Alexander, 63
evidence
for AI mental health tools, 186
evidence-adoption gap, 10
gap on patient outcomes, AI diagnostics, 91
integration of, 98–99
for wearables, 137–141, 152
Executive Order 14110, 66, 72, 285
Executive Order 14148, 285
Executive Order 14179, 285
explainability, AI, 13, 282, 287, 294
Explorys, 8
Eye2Gene project, 98
EyeArt, 114

F

Facebook, 179
fairness metric impossibility theorem, 69
fairness metrics, 291–292
false negatives, 118, 119, 180
false positives, 89–90, 118, 137, 145, 180

Fast Healthcare Interoperability Resources (FHIR) standard, 147, 152, 212, 215–216

Faves framework, 284

FDA. *See* U.S. Food and Drug Administration

federal AI legislation, 297–298, 321–322, 328

Federal Communications Commission, 14

Federal Food, Drug, and Cosmetic Act, 147

Federal Trade Commission (FTC), 149, 178–179

federated learning, 98, 114

Federation of State Medical Boards (FSMB), 93, 274, 286, 287, 293, 321

Fernihough, Donna, 63

FFR-CT. *See* fractional flow reserve from CT

FHIR. *See* Fast Healthcare Interoperability Resources standard

fiduciary duties, 250

Fitbit Heart Study, 140

Fitzpatrick skin tone scale, 57

Flo Health, 149

FLS. *See* Fundamentals of Laparoscopic Surgery

Food and Drug Omnibus Reform Act (2022), 15, 284

Force Feedback technology, da Vinci 5 (Model Is5000), 111

foundation models, 98, 114, 124–125, 139, 155, 222

fractional flow reserve from CT (FFR-CT), 210

Francisco Partners, 9

Freespira, 173

FreeStyle Libre, 140

FSMB. *See* Federation of State Medical Boards

FTC. *See* Federal Trade Commission

Fundamentals of Laparoscopic Surgery (FLS), 123

FURM framework, Stanford Health Care, 294, 297, 331

G

G42, 243

GAD. *See* generalized anxiety disorder

Galleri blood test, 84

Garcia, Megan, 177–178

Garcia v. Character Technologies, 177–178, 182

GDPR. *See* General Data Protection Regulation, European Union

gender bias, 59–60

General Data Protection Regulation (GDPR), European Union, 122, 179, 287

generalizability failures, 90, 118–119

generalized anxiety disorder (GAD), 171–172

generative AI, 67, 86, 169, 171–172, 176, 181, 182, 186, 187, 216, 217, 278, 283, 295, 316

Generative Pre-trained Transformer 4 (GPT-4), 59–60

Ghebreyesus, Tedros Adhanom, 250

GI Genius, 86, 116

Google, 137, 148, 178, 179, 182

  AMIE, 45, 222

  DeepMind, 21–22, 221, 279–280

  Med-PaLM 2, 86

  Pixel Watch 4, 138

Google Cloud, 96, 208

Google Health, 65, 88
Google Maps effect, 281
Google Research, 71
governance, 237, 240–241, 321–322
AI-assisted surgery and imaging, 122
AI diagnostics, 96–97
ambient AI documentation, 44–45
Assurance Standards Guide, CHAI, 70, 292
and autonomous AI systems, 296
data, 186–187, 217
infrastructure, 238
and innovation, gap between, 321
international frameworks, 250–251
and market trajectory, 259
philosophy, of Kaiser Permanente, 242
practices, third-party auditing of, 294
“Principles for AI Development, Deployment, and Use,” AMA, 251
RUAIH framework, 41, 97, 290, 292–295, 298, 326
state laws, 249
structure, 20
two-tier model, 254–255
Governance for Augmented Intelligence toolkit, AMA, 70, 93, 97, 251–252, 325–326
GPT-4. *See* Generative Pre-trained Transformer 4

H

hallucinations, AI, 37–38, 282

Hartree Centre, 209
Harvard Medical School, 325
Hassabis, Demis, 21
Hastings Center Report (2025), 181, 185
Hatchett, Sarah, 243
Hattiesburg Clinic, 43, 44
HCA Healthcare, 45, 85, 87–88
Health Breach Notification Rule, FTC, 149, 179
Health Canada, 214, 247, 284
healthcare, AI in, 2–3, 202, 239. *See also* medical devices, AI-enabled
  blind spots, 8–10
  clinical adoption, 4–5, 14
  ethical implications, 17–19, 216–218
  FDA clearances, 3–4
  future outlook, 21–23
  landmark successes, 5–8
  leadership implementation framework, 19–21
  market, 4, 205, 259
  patient perspectives on, 314–315
  real-world case studies, 10–11, 208–211
  regulatory and legal landscape, 15–17, 214–216, 283–286
  risks and limitations, 12–14, 244–247, 281–283
Health Care AI Accreditation program, URAC, 294

healthcare in 2030, AI in, 202–204
AI capability roadmap, 204
authoritative frameworks, 207–208
clinician adoption, 207
FDA authorization trajectory, 206
future outlook, 221–222
leadership implementation framework, 218–221
market and economic imperative, 205
reactive AI vs agentic AI, 204
risks and limitations, 211–214
workforce crisis, 205–206
Healthcare Leadership Council, 321–322, 328
health expenditures, projection of, 202, 205
Health Information Privacy Reform Act, 149
Health Tech Investment Act (2025), 215
Heartline study, 155
HHS (Department of Health and Human Services), 66, 211, 216, 283, 285, 293
HIPAA, 40, 121, 137, 149, 155, 179, 186, 282, 283, 327
HITECH Act, 30
HL7, 152
Hoberman, Brian, 35–36
Hospital AI Readiness Index, 208
Hospital at Home model, 220, 331–332
HTI-1. *See* ONC Health Data, Technology, and Interoperability Final Rule
HTNF. *See* Hypertension Notification Feature, Apple Watch

Hugo Ras, 111
Humana, 67
human-AI collaboration, 310–311
  ethical implications, 323–325
  evidence for, 311–315
  future outlook, 330–332
  leadership implementation framework, 325–330
  real-world case studies, 315–318
  regulatory and legal landscape, 321–322
  risks and limitations, 318–321
human oversight of AI mental health tools, 187
Hypertension Notification Feature (HTNF), Apple Watch, 138
hypoglycemia, 141
hypoxemia, occult, 12, 58–59, 145, 276

I

IBM, 2, 8–9, 61, 244–245
IBM Quantum System One, 209
Icahn School of Medicine, 180, 327
IDx-DR, 114

imaging, AI in, 5–6, 108–109, 125–126, 210. *See also* radiology, AI in
ethics, 121, 122
evidence base, 112–114
future outlook, 124–125
human-AI collaboration, 312
leadership implementation framework, 122–124
real-world case studies, 115–117
regulatory fragmentation, 120–121
risks and limitations, 118–120
incidental findings, diagnostic imaging, 94
incident reporting channels, 294–295
informed consent, 18, 40–42, 69, 121, 183, 188, 217, 279–280, 287, 293, 316, 324, 329
infrastructure, AI, 20
innovation champions, 44
Insilico Medicine, 221
insurers
use of wearable data, 150
wellness programs, 143–144
integration, AI, 19, 23, 63, 121, 187, 213
Intermountain Health, 40
Internal Review Committee (IRC), 254
International Medical Device Regulators Forum, 148
International Skin Imaging Collaboration (ISIC) dataset, 56, 276
Internet of Medical Things (IoMT), 207–208
interoperability, 152, 155, 215–216

Interoperable Telesurgery Protocol, 120
interstitial pulmonary fibrosis, progressive, 118
Intervention Risk Management (IRM), 248, 285
intestinal anastomosis, 112
Intuitive Surgical, 109–111, 125
investment, AI, 19–20, 237, 240
IoMT. *See* Internet of Medical Things
IRC. *See* Internal Review Committee
iRhythm Zio, 141
IRM. *See* Intervention Risk Management
Irregular Heart Rhythm Notification feature, Samsung, 138
Irregular Rhythm Notification Feature, Apple, 147, 148
ISIC. *See* International Skin Imaging Collaboration dataset
ISM001-055, 221
I-Thou relationship, 323

J

JAMA Health Policy Summit, 10
JAMA Network Open, 34–35, 38, 46, 174, 189, 258, 314
John Hancock, 143
Johns Hopkins, 7, 85, 88, 222, 317, 324, 325
Johnson & Johnson, 111
Joint Commission, 16, 22, 41, 44, 97, 218, 290, 292–294, 297, 326
Josiah Macy Jr. Foundation, 220, 327, 328
Jumper, John, 21

justice (ethics), 277–278

K

Kaiser Permanente, 10–11, 22, 29, 34–37, 42, 44, 96, 203, 209–210, 219, 221, 241–242, 258, 287–288, 315–316, 324

Katz, Mitchell, 243

Kheiron Medical Technologies, 87

Kotter’s eight-step model, 246

KPMG, 205

L

laboratory-developed tests (LDTs), 92

large language models (LLMs), 6, 15, 188, 246

- and bias, 12, 17, 184, 281
- and diagnostic partnership, 45–46
- and diagnostic reasoning, 86
- hallucinations, 282
- regulations, 17

LDTs. *See* laboratory-developed tests

leadership implementation framework, 19–21, 288, 290–295

- AI-assisted surgery and imaging, 122–124
- AI mental health tools, 185–188
- ambient AI documentation, 43–45
- clinical AI deployment, 70–71
- human-AI collaboration, 325–330
- predictive diagnostics, 95–98
- wearables, 151–153

lead-time bias, 89

left ventricular systolic dysfunction, 154
Lenox Hill Hospital, 125
LGBTQ+ individuals, and AI mental health tools, 183, 184
liability, 93, 216–218, 283, 287, 296
  AI liability chain, 290
  AI mental health tools, 182–183, 188
  framework, lack of, 17
Life Sciences and Healthcare Predictions 2030 series, Deloitte, 207–208
Limbic, NHS, 175
literacy, AI, 218, 319, 325, 327–328
LLMs. *See* large language models
Loris.ai, 179
Loss of Pulse Detection feature, Google, 138, 148
Lozovatsky, Margaret, 326
Lumeris, 222
LumineticsCore, 217
lung cancer, prediction of, 84
lung nodules, detection of, 116
Lunit Insight Cxr, 113
Lyra Health, 176, 187

M

machine learning, 7, 60, 63, 139, 282
Mai-DxO, 222
major depressive disorder (MDD), 171–172, 189

malpractice, 216, 287
and AI mental health tools, 188
and ambient AI documentation, 41
involving AI diagnostic tools, 91, 93
mammography, 5, 46, 83, 87, 89–91, 95, 108, 109, 113, 118–119, 126, 210, 312
Marchand v. Barnhill, 250
MASAI (Mammography Screening with Artificial Intelligence) trial, 5, 46, 83, 89, 109, 113, 126, 312
Mass General Brigham, 34, 143, 152, 258, 313, 331
Mater Hospital Dublin, 222
Maven Clinic, 144
Mayo Clinic, 84, 91, 96, 143, 153, 203, 208–209
Mayo Clinic Medical School, 218
Mayo Clinic Platform, 208–209
MCED. *See* multicancer early detection tests
McKinsey & Company, 205, 239, 258, 260
MD Anderson Cancer Center, 61, 244
MDD. *See* major depressive disorder
Meaningful Use requirements, 30
Medicaid, 119, 151
Medical Device Regulation (EU 2017/745), European Union, 148–149, 286

medical devices, AI-enabled, 60, 63, 67, 118, 237, 240–241, 274, 275
bias reporting for, 71
EU AI Act compliance requirements, 214–215
FDA draft guidance on AI-Enabled Device Software Functions, 72, 284
FDA regulation, 3–4, 15, 66, 85, 92, 108, 112, 125, 206, 214, 284
growth of, 2, 3
recalls, 10, 67, 214, 247
and wellness apps, boundary between, 184–185
Medical Devices Regulations (2002), United Kingdom, 148
medical education, 218, 325, 327–328
medical errors, and physician burnout, 29, 32–33
medical humanities, 324
medical supply chain, 317
Medicare, 93, 114, 181, 205, 322, 331
Advantage, 67, 144, 153, 155, 251
New Technology Add-on Payment, 115
RPM programs, 151
Shared Savings Program, 219
Medicines and Healthcare products Regulatory Agency (MHRA), United Kingdom, 148, 214, 247, 284, 322
Med-PaLM 2, 86
MedSAM, 98, 114, 124
MedSimAI, 328
Medtronic, 86, 111, 116
melanoma, 209, 276
Memorial Sloan Kettering Cancer Center, 9, 61

Menlo Ventures, 209, 240
mental health, 168
services, integration of AI mental health tools with, 187
treatment gap, 169–171
mental health, AI in, 168–170
chatbots, 169, 171–173, 175–176, 179, 182, 183, 185–188
ethical implications, 183–185
future outlook, 188–190
leadership implementation framework, 185–188
market, 171
prescription digital therapeutics, 173–174
real-world cases, 174–176
regulatory and legal landscape, 180–183
risks and limitations, 177–180
safety spectrum, 177
Mental Health Network, NHS Confederation, 175
Merge Healthcare, 8
Meta Segment Anything Model 2, 125
MFMU Network, 213
MHRA. *See* Medicines and Healthcare products Regulatory Agency, United Kingdom
Mia (breast cancer screening AI), 87
Microsoft, 6, 10, 33, 40, 222
Microsoft Research, 114, 222
Mihaljevic, Tom, 243
Miller, Chris, 243–244

MIMIC-III database, 60
MIMIC-IV database, 59
Minimally Invasive Robotic Association (MIRA), 123
MIRA. *See* Minimally Invasive Robotic Association
Mishuris, Rebecca, 34
Mobile trial, 141
Model Card Registry, CHAI, 41, 44
model cards, 71, 292, 331
model drift problem, 181
monitoring
  of AI mental health tools, 187
  algorithm vigilance, 96, 122
  ambient AI documentation, 45
  continuous, and wearables, 150
  post-deployment, 10, 17, 45, 72, 96–98, 122, 123, 214, 241, 247
Monk Skin Tone Scale, 145
MoodCapture, 189
mortality prediction systems, 296
mosaic effect, 121–122
multicancer early detection (MCED) tests, 84, 95, 99
multimodal AI systems, 23, 98–99, 125, 189
My Intuitive+ analytics ecosystem, 111

N

Nabla, 33, 38
NAM. *See* National Academy of Medicine
narrative medicine, 323–324

National Academy of Medicine (NAM), 38, 70, 298, 324, 329–330
National Eating Disorders Association (NEDA), 175–176, 183
National Health Service (NHS), United Kingdom, 175, 203, 211
  Foresight initiative, 210
  NHS Confederation, 175
  NHS-Galleri trial, 84, 99
  NHS Grampian breast screening pilot, 87
  pilot purgatory, 213
  Royal Free London NHS Foundation Trust, 279–280
National Kidney Foundation, 69
National Lung Screening Trial (NLST), 84, 89
natural language processing (NLP), 117, 171, 210
Nature, 45
NCSBN, 31
Nebraska Medicine, 219
NEDA. *See* National Eating Disorders Association
negative skill transfer, 119
NEJM AI, 6, 33, 70
networked medical devices, 120
neurofeedback, AI, 189
never-skilling, 319, 328
New Technology Add-on Payment (NTAP), Medicare, 115
next-generation wearables, 154
nH Predict algorithm, 67, 280
NHS. *See* National Health Service, United Kingdom
NIH, 72, 99

NIST, 72, 253

NLP. *See* natural language processing

NLST. *See* National Lung Screening Trial

Noni, 178

nonmaleficence (ethics), 277–278

Northwell Health, 125

Northwestern Medicine, 151, 152, 219

“Now, Next, and Always” framework, 238, 260

NTAP. *See* New Technology Add-on Payment, Medicare

Nuance, 6, 221

nurse burnout, 31, 33

nursing shortage, 206

NYC Health + Hospitals, 243

O

obstructive sleep apnea, 208

occult hypoxemia, 12, 58–59, 145, 276

OECD. *See* Organisation for Economic Cooperation and Development

Omicron wave, COVID-19 pandemic, 30

omission errors, 281

ONC Health Data, Technology, and Interoperability Final Rule (HTI-1), 248, 284–285

one-party consent states, 40

OpenAI, 178

ophthalmology, AI in, 8, 114

Optum, 55, 56, 245, 252, 274–276, 278–279

Organisation for Economic Cooperation and Development (OECD), 251

orthosomnia, 146
Ottava, 111
Oura Health, 144, 154
Oura Ring, 139, 142, 144
outcome measurement for AI mental health tools, 187
overdiagnosis, 89

P

Paige, 86, 113
pajama time, 29, 32, 313, 323
Palantir Technologies, 209, 243
palliative care, 296
PANDA AI, 46
Pathak, Divya, 243
Pathfinder 2, 99
pathology, AI in, 8, 86, 113–114
patient autonomy, 18, 94, 122, 277–278, 287
patient consent. *See* informed consent
patient notification protocols, 293
patient perspectives on healthcare AI, 314–315
Patient Safety Organizations (PSOs), 295
pattern recognition, 8, 56, 86, 87
Payne, Philip R.O., 243
PCCPs. *See* Predetermined Change Control Plans, FDA
PDSA. *See* Plan-Do-Study-Act cycle
PDTs. *See* prescription digital therapeutics
Pearl-AI. *See* Pediatrics EthicAl Recommendations List for AI framework

Pear Therapeutics, 173–174, 181

pediatric AI ethics, 296–297

Pediatrics EthicAl Recommendations List for AI (Pearl-AI) framework, 296–297

Penda Health/OpenAI AI clinical copilot study, 217, 222

People, Process, Technology, and Operations (PPTO) framework, 254–255

Peralta, Juliana, 178

performance paradox, 14

performance reviews, 97

Permanente Medical Group, The (TPMG), 10–11, 22, 34–37, 44, 45, 218, 258, 313

phased implementation, healthcare AI, 220–221, 256–257

PHI. *See* protected health information

photoplethysmography (PPG), 138, 139, 144–145

physician burnout, 143, 206, 219, 258, 288, 313, 314. *See also* ambient AI documentation

- costs of, 31
- data on, 30–32
- and documentation burden, 6–7, 28–29, 32–33
- rate, tends of, 35
- systemic nature of, 38

physicians. *See also* clinician adoption; human-AI collaboration

- accountability, 286, 321
- AI literacy of, 325
- empathy of, 315, 323
- physician champions, 21, 44, 253–254
- physician-patient relationship, 18–19, 39, 218, 287–288, 323, 324

role in AI governance, 325–326
shortage, 31, 202, 203, 205–206
well-being of, 42, 43
Physicians Foundation, 31, 38
Phytel, 8
pilot(s), 257
for ambient AI documentation, 44
purgatory, 213
Pixel Watch 4, 138
Plan-Do-Study-Act (PDSA) cycle, 44
Plasmodium falciparum malaria, 117
Platform Orchestrate, Mayo Clinic, 209
post-deployment monitoring, 10, 17, 45, 72, 96–98, 122, 123, 214, 241, 247
PPG. *See* photoplethysmography
PPTO. *See* People, Process, Technology, and Operations framework
PRAIM trial, 83
Predetermined Change Control Plans (PCCPs), FDA, 4, 15, 22, 66, 92, 113, 148, 155, 181, 206, 214, 241, 247, 284
predictive AI, 4, 236, 238, 326

predictive diagnostics, 80–81
AI-assisted radiology workflow, 81
cancer detection, 83–84
cardiovascular disease, 84–85
ethical implications, 93–95
future outlook, 98–99
leadership implementation framework, 95–98
real-world case studies, 86–88
regulatory and legal landscape, 91–93
risks and limitations, 89–91
sepsis, 85
prescription digital therapeutics (PDTs), 173–174
pressure ulcer severity, analysis of, 211
“Principles for AI Development, Deployment, and Use,” AMA, 251
privacy, 139, 149, 152
re-identification of patients, 282–283
surgical video, 121–122
violations, and AI mental health platforms, 178–179
proactive ethics, 297
“problem-first, technology-second” approach, 238
productivity trap, and ambient AI documentation, 42, 43
Proprio, 125

prostate cancer

detection of, 86, 113

grading, 8

protected health information (PHI), 121–122, 216, 245–246, 279–280, 282–283

protein structure prediction problem, 21–22

Providence Health, 6, 34, 42, 43

PSOs. *See* Patient Safety Organizations

Public Health AI Handbook, 255

pulse oximetry, 12, 58–59, 145–146, 276–277

PwC, 208

Pxdx algorithm, 67

Q

quality-adjusted life years (QALYs), 91, 153

Qure.ai, 113, 117

qXR, 113, 117

qXR-Detect, 113

R

racial bias, 2–3, 12, 54, 69, 212–213, 320
in AI diagnostics, 94
in AI imaging, 120
in AI mental health tools, 183–184
DeepDerm, 57, 276
Optum, 56, 245, 252, 274–276, 278–279
pulse oximetry, 58–59, 276–277
racial disparities, 217
and AI adoption gap, 320
in AI false positives, 89–90
health disparities, 150
RACOON initiative, 114
Radiological Society of North America (RSNA), 64
radiology, AI in, 5–6, 83, 85–86, 90, 92, 108, 312. *See also* imaging, AI in
adoption, 114
and disclosure, 121
radiologist-AI collaboration, 114
ROI measurement, 258
workflow, 81
Raine, Adam, 178
Raine v. OpenAI, 178, 182
Ralph, Erin, 63
ransomware attack, 211–212, 245–246
rapid response protocol, for AI-related incidents, 295
Raven II, 120
RCTs, ambient AI documentation, 33–34, 38–39

REACH VET, 174–175, 316

reactive ethics, 297

readmission reduction program, AI-integrated, 42–43

REAL trial, 109, 111, 126

reasonable physician standard, 216

recalls, AI-enabled medical device, 10, 67, 214, 247

“Recommendation on the Ethics of Artificial Intelligence,” UNESCO, 277

Recovery Engagement and Coordination for Health: Veteran Enhanced Treatment (REACH VET), 174–175

rectal cancer surgery, 109, 111, 126

ReDS Wearable System, 141

regulations, 13, 15–17, 214–216, 240–241, 283–286

AI mental health tools, 180–183

and ambient AI documentation, 40–41

clinical AI deployment, 65–68

gaps, 17

human-AI collaboration, 321–322

imaging, 120–121

maturation of, 22

predictive diagnostics, 91–93

wearables, 147–149

re-identification of patients, 282–283

reimbursement, 93, 114, 115, 151, 212, 215, 249, 321, 322, 328

Rejoyn, 173

remote patient monitoring (RPM), 143, 151, 153, 220. *See also* wearables

Remote Patient Monitoring–Home Telehealth program (RPM-HT), VA, 142–143, 151

reporting
bias, 71
incident reporting channels, 294–295
transparency in, 45
Responsible Use of AI in Healthcare (RUAIH) framework, 41, 97, 290, 292–295, 298, 326
right not to know, 93–94
right to know, 329
right to refuse AI, 122
Rigshospitalet, Copenhagen, 61
risk assessment of AI mental health tools, 185–186
robotic surgery, 108–110, 123, 126
adverse events, 117–118
autonomous, 112, 124
cost of surgical robots, 119
ROI
for AI imaging, 123
for ambient AI documentation, 219
measurement of, 20, 258–259
for wearables, 153
ROLARR trial, 112
Royal Free London NHS Foundation Trust, 279–280
RPM. *See* remote patient monitoring
RPM-HT. *See* Remote Patient Monitoring–Home Telehealth program, VA
RSNA. *See* Radiological Society of North America
RUAIH. *See* Responsible Use of AI in Healthcare framework

rural healthcare AI gap, 14, 212, 216, 320

Rural Health Transformation Program, 214

S

saddle pulmonary embolism, 142

safety protocols, for AI mental health tools, 186

SAGES. *See* Society of American Gastrointestinal and Endoscopic Surgeons

SaMD. *See* Software as a Medical Device framework

Samsung, 137, 138, 147, 154

Samsung Galaxy Watch 6, 142

Samsung Health Monitor app, 138

Samsung Medical Center (SMC), 203, 210–211

Schabacker, Marcus, 252

Scheel, Paul, 243

Scribe, 243

sepsis

  Epic Sepsis Model, 9, 13, 21, 62, 90, 245

  prediction of, 2, 7, 85, 87–88

  Sepsis Watch, 85, 96

Setzer, Sewell, III, 177

Shahshahani, Ben, 242

Shanafelt, Tait, 30, 44

shortcut learning, 91

SimBioSys, 125

sinuplasty, 63

skin cancer diagnosis, 56–58, 276

Skinex AI, 211
skin lesion analysis tools, 8
skin tone
and AI dermatology models, 69
Fitzpatrick skin tone scale, 57
and SpO_2 measurement, 145–146
sleep apnea, 138, 139, 147, 208
SleepioRx, 173
smart biosensor patches, 141
Smart Tissue Autonomous Robot (STAR), 112
SMC. *See* Samsung Medical Center
Smitherman, Scott, 43
Snapchat, 179
Society of American Gastrointestinal and Endoscopic Surgeons (SAGES), 123
Software as a Medical Device (SaMD) framework, 148
SpO_2, 145–146
Spot (Sepsis Prediction and Optimization of Therapy) system, 85, 87–88
Spring Health, 176, 186, 187
SRT-H. *See* Surgical Robot Transformer-Hierarchical
standard of care, 22, 46–47, 93, 188, 283
Standing Together consensus recommendations, 70–71
Stanford Health Care, 294, 297, 331
Stanford Professional Fulfillment Index, 34
Stanford School of Medicine, 32, 38–39, 44, 56, 218, 325, 327
STAR. *See* Smart Tissue Autonomous Robot

state laws, 181–182, 249, 285, 297, 298
statistical parity (fairness metric), 291
Stelo, 141
Storm (Stratification Tool for Opioid Risk Mitigation), 316
Streams, 279–280
stroke imaging, 5–6, 108, 115
substance use disorder, 172
suicide, 170, 180
  and chatbots, 177–178
  prediction, 174–175, 184
Suki (ambient AI), 37
surgery, AI in, 108–109, 125–126
  ethics, 121–122
  evidence base for, 110–112
  future outlook, 124, 125
  leadership implementation framework, 122–124
  risks and limitations, 117–120
surgical coaching, AI, 125
surgical robotics. *See* robotic surgery
Surgical Robot Transformer-Hierarchical (SRT-H), 112, 124
surgical video privacy, 121–122
Sutter Health, 38, 45, 59, 116
Sybil model, 84
Synapsis AI, 209

T

talent shortages, 246

TEFCA. *See* Trusted Exchange Framework and Common Agreement
telehealth use, disparities in, 288, 289
telemedicine, 318
Tempo pilot, 155
TemPredict study, 142
Tessa, 175–176, 183
Therabot, 169, 171–172, 187, 190
therapeutic relationship, 18–19, 39, 218, 287–288, 323, 324
third-party auditing of governance practices, 294
three-role framework, 252–254
TikTok, 179
Total Product Lifecycle (TPLC) approach, 247
TPLC. *See* Total Product Lifecycle approach
TPMG. *See* Permanente Medical Group, The
training, 244
  on AI-assisted surgery and imaging, 123–124
  on AI diagnostic tools, 95
  on AI mental health tools, 187
  building AI workforce, 220
  for human-AI collaboration, 327–328
  programs, for ambient AI documentation, 44
training data, 54–55, 320
  cardiovascular AI, 59–60
  dataset shift, 65
  DeepDerm, 57, 276
  Epic Sepsis Model, 62

and equity, 18
geographic concentration of, 60
Model Cards and Datasheets for Datasets, 71
Optum algorithm, 56, 252, 274–276, 278–279
preprocessing interventions, 291
quality of, 13, 60
racial/ethnic composition of, 12, 17, 60, 184
representation of low-income population in, 54
representative, 72
Watson for Oncology, 8–9, 61
Transforming Episode Accountability Model (Team), CMS, 219
Transitional Coverage for Emerging Technologies pathway, 93
transparency, 9, 19, 64, 95, 122, 282, 324, 329
and AI mental health tools, 183
and ambient AI documentation, 41–42, 45
attributes, 248, 284–285
of clinical AI studies, 12
guiding principles for ML-enabled devices, 13
Model Cards and Datasheets for Datasets, 71
TREWS (Targeted Real-Time Early Warning System), 7, 85, 88
TruDi Navigation System, 63, 118
Trump, Donald, 66, 72, 284, 285

trust
clinician trust deficit, 14, 213
gap, 203
between patient and physician, 19, 314
of patients in AI, 314–315, 319–320, 329
and transparency, 45
Trusted Exchange Framework and Common Agreement (TEFCA), 215
Truven Health Analytics, 8
tuberculosis, detection of, 117
TumorSight Viz, 125
two-party consent states, 40

U

UCLA Health, 33, 37, 38, 40, 45
UCSF Algorithmic Justice Workshop, 68, 71
UCSF Health, 96
UNESCO, 277
unified clinical intelligence, 45
UnitedHealthcare
nH Predict, 67, 280
Rewards program, 144
University of Chicago Medicine, 314
UPMC, 151
URAC, 294

U.S. Food and Drug Administration (FDA), 2–5, 8, 17, 58, 66, 72, 92, 97, 108, 113–114, 116, 121, 122, 145, 237, 284, 285, 328
AI/ML-Based Software as a Medical Device Action Plan, 15
and ambient AI scribes, 40
authorization trajectory, 206, 240–241
clearance for surgical robots, 111
digital therapeutics cleared by, 173–174
framework for AI mental health tools, 180–181, 186
LDT regulation, 92
MAUDE data, 117
post-market monitoring infrastructure, 214
Predetermined Change Control Plans, 4, 15, 22, 66, 92, 113, 148, 155, 181, 206, 214, 241, 247, 284
regulation of wearables, 137, 147, 152, 155
Software as a Medical Device framework, 148
Total Product Lifecycle approach, 247
transparency guiding principles for ML-enabled devices, 13
UW Health, 33–34, 38, 44, 45

V

VA. *See* Department of Veterans Affairs
vaginal birth after cesarean (VBAC), 213, 217
VA GPT, 316
VALIDATE study, 115
validation of AI tools, 21, 65
value-based care, 219, 260, 322
Vanderbilt Suicide Attempt and Ideation Likelihood (VSAIL) model, 189

VBAC. *See* vaginal birth after cesarean

vendor(s)

accountability, in contracts, 294

evaluation criteria, 152, 186–187

vendor evaluation scorecard, 255–256

Vera-MH framework, 186

Versius Plus, 111

Vertex AI, 208

Veterans Administration. *See* Department of Veterans Affairs (VA)

Veterans Health Administration, 142–143, 174–177

Virtual Command Center, Cleveland Clinic, 209, 243

virtual reality therapy, AI-guided, 189

VitalPatch, 141

vital signs, EHR, 60

Viz.ai, 6, 108, 115

Viz LVO, 115

voice-based depression screening tools, 189

VSAIL. *See* Vanderbilt Suicide Attempt and Ideation Likelihood model

vulnerable populations, protection of, 185

W

Wachter, Robert, 38

Washington My Health My Data Act (2024), 149

Washington University School of Medicine, 243–244

Wasteful and Inappropriate Service Reduction (WISeR) model, 249, 260

Watson for Oncology, 8–9, 61, 244–245

Watson Health, 2, 8–9, 61

wearables, 136–137, 220
digital phenotyping using, 189
ethical implications, 149–150
evidence base for, 137–141, 152
future outlook, 154–155
health data, 146–147, 149
landscape of, 137–138
leadership implementation framework, 151–153
real-world case studies, 141–144
regulatory and legal landscape, 147–149
risks and limitations of, 144–147
well-being, physician, 42, 43
Wellness and Oversight for Psychological Resources (WOPR) Act, Illinois, 181–183, 188
White House Blueprint for an AI Bill of Rights, 251
WHO. *See* World Health Organization
whole-slide imaging analysis, AI-assisted, 8
WISeR. *See* Wasteful and Inappropriate Service Reduction model
Witty, Andrew, 245
Woebot, 172
Woebot Health, 180
WOPR. *See* Wellness and Oversight for Psychological Resources Act, Illinois
workflow redesign, 21

workforce
AI, building, 220
crisis, 31, 202, 203, 205–206, 246
development, 246
disruption, 320–321
transition funding, 328
World Economic Forum, 208
World Health Organization (WHO), 16, 60, 68, 95, 117, 155, 212, 217, 250, 277, 295, 324
Worthington, Charles, 317
Wysa, 172

Y

Yang, Daniel, 242

Z

Zipline, 317
ZS Associates, 321–322
Zuckerberg San Francisco General Hospital, 42–43

AI in Emergency Medicine: What All Healthcare Professionals Can Learn from the Artificial Intelligence Revolution, First Edition. Harvey Castro.
© 2027 by John Wiley & Sons Inc. All rights reserved, including rights for text and datamining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Copyright © 2027 by John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permission.

The manufacturer's authorized representative according to the EU General Product Safety Regulation is Wiley-VCH GmbH, Boschstr. 12, 69469 Weinheim, Germany, e-mail: Product_Safety@wiley.com.

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

Limit of Liability/Disclaimer of Warranty: While the publisher and the author have used their best efforts in preparing this work, including a review of the content of the work, neither the publisher nor the author make any representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives, written sales materials or promotional statements for this work. The fact that an organization, website, or product is referred to in this work as a citation and/or potential source of further information does not mean that the publisher and author endorse the information or services the organization, website, or product may provide or recommendations it may make. This work is sold with the understanding that the publisher is not engaged in rendering professional services. The advice and strategies contained herein may not be suitable for your situation. You should consult with a specialist where appropriate. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002. For product technical support, you can find answers to frequently asked questions or reach us via live chat at https://support.wiley.com.

If you believe you've found a mistake in this book, please bring it to our attention by emailing our reader support team at wileysupport@wiley.com with the subject line "Possible Book Errata Submission."

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Control Number: 2025950800

ISBN 9781394385096 (Paperback)
ISBN 9781394385102 (ePub)
ISBN 9781394385119 (ePDF)

Cover Design: Wiley
Cover Image: © metamorworks/Getty Images

To my mother, who arrived in this country at 13 with nothing and gave me everything. And to my children, who remind me every day why the future of healthcare matters.

Acknowledgments

This book would not exist without the people who supported it at every stage.

Thank you to my editor, Cathleen Small, whose sharp eye and thoughtful questions made every chapter stronger. To Satish and the entire Wiley team for believing in this project and guiding it to the finish line.

To my mother, who taught me that showing up for people is never optional. To my children, who tolerated the late nights and weekend writing sessions with patience I probably did not deserve.

And to every clinician, administrator, and patient who trusted me with their time and their stories. This book is built on what you taught me.

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.

This file was downloaded from Z-Library project



Your gateway to knowledge and culture. Accessible for everyone.

z-library.sk z-lib.gs z-lib.fm go-to-library.sk

Official Telegram channel

Z-Access

https://wikipedia.org/wiki/Z-Library