

Driving Adoption of High-Value AI in Healthcare

MARKET TREND REPORT | OCTOBER 2026




Table of Contents

2	Letter from the Executive Director	→
4	Expert Contributors	→
7	Executive Summary	→
10	About This Report	→
11	An Opportunity to Transform Healthcare: Why AI Is Unique	→
14	Accelerating Broken Processes: The Risks of Administrative AI	→
17	Unrealized Potential: The Promise of AI for Clinical Care	→
21	Setting the Right Incentives: How Payment Models Drive AI Adoption	→
25	The Path Forward: Proactive Development of High-Value AI	→
29	Paying for AI in Healthcare: Three Principles to Drive Value	→
30	Looking Ahead	→
31	References	→

Letter from the Executive Director

As the busy mother of two school-aged kids, I know how stressful it is to not be able to get them the care they need quickly and reliably. There are monthly struggles with medication refills, not knowing whether they need to see a doctor after office hours, or finding high-quality clinical care for more complex conditions. Even as a healthcare expert, I regularly bump up against the failures of our system. For families across the country, especially those in rural and underserved communities, today's healthcare is too often hard to navigate, slow to access, infuriatingly opaque, and prohibitively expensive. There simply isn't enough clinician attention to go around, and this is part of why costs keep climbing and why people who need care don't get it.

This report, which calls for accelerated development and deployment of clinical applications of AI, comes from a sober perspective: the road ahead will take a lot of work to get right, and we can't afford not to try. We will take safety seriously. We will monitor performance and impact along the way. And we will remain attuned to real concerns from patients and clinicians.

But let us be clear: If AI can improve even a fraction of today's healthcare challenges, we have to do the hard work to realize its potential. Not out of a cavalier belief in the power of technology, but out of the urgent need to deliver better, more accessible and affordable care to our families.

Take one example: high blood pressure. Four out of five people diagnosed with hypertension are not in blood pressure control. Most of those people see a doctor regularly, and we have low-cost medications that are effective at treating hypertension. Yet, amidst many other demands on a primary care doctor's time, too often high blood pressure is not the top priority in a given visit. But what if AI-enabled technology could monitor and manage blood pressure between visits—getting more patients into control and leaving time for doctors to focus on other issues?

AI already assists clinicians with existing tasks, including clinical decision support, documentation, and billing. Increasingly, it can take on aspects of direct care delivery, such as screening and diagnosis, remote monitoring, and medication prescribing. That means patients can receive care between visits, and clinicians can focus on the more complex cases that demand their expertise and attention.

In considering whether to adopt AI in clinical care, we cannot compare its impact against some idealized, always-available physician with endless knowledge and patience. We need to ask whether the solution can improve outcomes in the real world that exists today. For too many patients, their reality is low-quality care, delayed care, or no care at all—due to cost, access, or both. Measured against that reality, AI-enabled care can offer real benefits if it is well tested, carefully deployed and overseen, and appropriately reimbursed.

Today, the clinical applications of AI with the greatest potential to expand access are stalled while administrative billing and documentation tools are proliferating and being used to drive up healthcare costs without improving health outcomes. On our current trajectory, AI is set to exacerbate existing system failures.

If, instead, AI is going to play a role in solving healthcare's most intractable problems, we need a deliberate and concerted effort to build the right incentives, guardrails, and delivery models. This demands a focused effort across four domains: (1) generate better evidence about what works; (2) enable coverage for AI tools that deliver care with limited clinician involvement by updating licensure and liability rules; (3) build payment models that reward outcomes rather than activity; and (4) redesign workflows and train the workforce to use AI effectively in care delivery.

PHTI will continue to advance these aims through partnership, research, and dialogue. Others in the sector must take action and lead where they have expertise and authority: policymakers need to modernize regulatory oversight and licensure to ensure safety and enable the use of effective solutions; health systems and providers will need to redesign care delivery models and support their workforce through this transition; and health technology companies and investors must build the evidence that proves what works.

AI offers us an opportunity to build the healthcare system that works for patients—affordable care that is accessible when we need it. This roadmap lays out the necessary steps that can make it a reality.

Sincerely,

A handwritten signature in black ink, reading "Caroline Pearson". The signature is fluid and cursive, with a long horizontal line extending from the end of the name.

Caroline Pearson, Executive Director
Peterson Health Technology Institute

Expert Contributors

This report was informed by a series of interviews and expert workshops from late 2025 through 2026 that included leaders from health systems, medical societies, health plans, technology companies, academia, investment firms, patient groups, and federal government agencies to identify the required market and policy conditions to scale high-value AI adoption across healthcare. We are grateful to the individuals listed below for generously sharing their time, expertise, and insights throughout these discussions. Their contributions helped deepen our understanding of the issues, challenge our thinking, and inform the conclusions reflected in this report. Inclusion as a contributor does not indicate endorsement of the report's findings or recommendations.

Peterson Health Technology Institute

- Caroline Pearson
Executive Director
- Prabhjot Singh, MD, PhD
Senior Advisor for Strategic Initiatives
- Corinna Cline
Manager, Assessments
- Julia Harris, MPH, MA
Senior Director, Policy
- Steven D. Pearson, MD, MSc
Senior Advisor

Contributors

- Rick Abramson, MD, MHCDS, FACR
Director, Digital Health Center of Excellence, U.S. Food and Drug Administration
- Matt Adkins
Senior Vice President, Strategy & Transformation, CenterWell Senior Primary Care
- Jerome Avondo, PhD
Vice President, Clinical & Reimbursement, Aidoc
- Suresh Balu
Associate Dean for Innovation and Partnership, Duke University
- Luke Belz
Vice President of Health Economics, Sword Health
- Ami Bhatt, MD, FACC
Chief Medical Officer, WHOOP
- Melinda Buntin, PhD
Bloomberg Distinguished Professor of Health Policy and Economics, Johns Hopkins University

- Sebastian Caliri
Partner, 8VC
- Michael Chernew, PhD
Professor, Harvard Medical School, Health Care Policy
- Richard Clarke, PhD
Chief Analytics Officer, Highmark Health (Former)
- Mandy Cohen, MD, MPH
Director, U.S. Centers for Disease Control and Prevention (Former); Secretary, North Carolina Department of Health and Human Services (Former)
- Byron Crowe, MD, MS
Chief Medical Officer, Doctronic; Clinical Assistant Professor, Stanford University School of Medicine
- Eve Cunningham, MD, MBA
Chief Medical Officer, Cadence
- Natalie Davis
Chief Executive Officer and Co-founder, United States of Care
- Niyum Gandhi
Chief Financial Officer and Treasurer, Mass General Brigham
- Vivek Garg, MD, MBA
Chief Medical Officer, CenterWell Senior Primary Care (Former)
- William Gordon, MD, MBI
Chief Informatics Officer, Centers for Medicare & Medicaid Services Innovation Center (Former)
- Elliott Green
Chief Executive Officer and Co-founder, Dandelion Health
- Duncan Greenberg
Venture Partner, Oak HC/FT

- Ajeet Hansra
Director, New Mountain Capital
- Brian Hasselfeld, MD
Executive Medical Director, Digital Health and Innovation, Johns Hopkins Medicine
- Zach Hochstetler
Vice President, CPT Coding and Payment, American Medical Association
- Tom Insel, MD
Co-founder, Benchmark Health and Vanna Health; Director, National Institute of Mental Health, National Institutes of Health (Former)
- Ankit Jain
Co-founder and Chief Executive Officer, Infinitus Systems
- Karen Joynt Maddox, MD, MPH
Professor of Medicine/Cardiology, Washington University School of Medicine and School of Public Health
- Goutham Kandru
Associate Director, U.S. Healthcare, Health and Life Sciences, Gates Ventures
- Chris Klomp
Chief Counselor, U.S. Department of Health and Human Services
- Adam Landman, MD
Chief Digital Information Officer, Brown University Health
- Joyce Lau
Director of Strategy and Technology Policy, Center for Medicare, Centers for Medicare & Medicaid Services
- Puneet Maheshwari
Senior Vice President and General Manager, Optum Real at Optum Insight
- Tarun Mehra
Partner, Healthcare & Life Sciences, Microsoft Corporation
- Joel Meyer
Public Sector, Anthropic
- David Mohr, PhD
Professor, Director of the Center for Behavioral Intervention Technologies, Northwestern University, Feinberg School of Medicine
- Amol Navathe, MD, PhD
Professor, University of Pennsylvania
- Shantanu Nundy, MD, MBA
Advisor on AI, Office of the Commissioner, U.S. Food and Drug Administration (Former)
- Neil Parikh, MD, MPH, MSM, FABQAURP, CPE
National Chief Medical Officer, Medical Management, UnitedHealthcare Clinical Services
- Bakul Patel
Senior Director, Global Digital Health Regulatory Strategy, Google
- Kavita Patel, MD, MS
Professor, Stanford University School of Medicine
- Philip R.O. Payne, PhD, FACMI, FAMIA, FAIMBE, FIAHSI
Chief Health AI Officer, BJC HealthCare and WashU Medicine
- Jacob Quinton, MD, MS, MPH
Chief Medical Officer, Medicare Parts A/B, Centers for Medicare & Medicaid Services
- Sandhya Rao, MD
Chief Medical Officer and Senior Vice President, Health and Medical Management, Blue Cross Blue Shield of Massachusetts
- Raj Ratwani, PhD, MPH
Vice President, Scientific Affairs, MedStar Health
- Anindita (Annie) Saha
Associate Director for Strategic Initiatives, Digital Health Center of Excellence, U.S. Food and Drug Administration (Former)
- Kevin Schulman, MD, MBA
Interim Co-Chief, Division of Hospital Medicine, Stanford University School of Medicine
- Mark Sendak, MD, MPP
Co-founder and Chief Executive Officer, Vega Health
- Jacob Shiff, MBA
Chief AI and Technology Officer, Centers for Medicare & Medicaid Services Innovation Center
- Benjamin Shobert
Chief Operating Officer, UnitedHealthcare Clinical AI
- Haider Warraich, MD
Program Manager, Scalable Solutions, Advanced Research Projects Agency for Health
- Jeff Wessler, MD
Founder and Chief Executive Officer, Heartbeat Health
- Daniel Yang, MD
Vice President of AI and Emerging Technologies, Kaiser Permanente
- Gigi Yuen, PhD
Chief Data and AI Officer, Cohere Health

Manatt Health provided facilitation and research support for the series of workshops and throughout the development of this report.

- Jared Augenstein, MPH, MA
Senior Managing Director
- Allie Levy Chafetz, MBA
Director
- Lili Muskal, MPH
Consultant
- Naomi Newman, MBA, MPH
Managing Director
- Joshua Tauber, MBA
Managing Director

The Peterson Health Technology Institute (PHTI) accelerates the adoption of effective, high-value health technologies to improve health, expand access, and lower spending. Our evidence-based research informs purchasing decisions by employers, providers, and health plans to ensure that digital health solutions deliver meaningful clinical benefits for patients. Founded in 2023 and funded solely by the Peterson Center on Healthcare, PHTI's research findings are independent with no financial ties to the companies or technologies assessed. Learn more at PHTI.org.

Executive Summary

Artificial intelligence (AI) has the potential to transform healthcare delivery and administration by performing clinical and administrative tasks that have previously been constrained by human expertise and workforce capacity. AI creates new capacity in two ways: it can enable the existing workforce to do more (assistive AI), and it can directly execute tasks and deliver care (autonomous AI). **If deployed effectively, AI could expand access to care, improve outcomes, and lower spending.** Savings generated by health AI should be used to reduce total system costs, including lower premiums and cost-sharing, for patients, employers, and taxpayers.

Whether AI realizes this potential, however, will depend heavily on the financial incentives that guide its deployment. When incentives reward transactions, documentation, billing activity, and administrative processes, AI accelerates those activities. If incentives reward improved patient outcomes, expanded access, and administrative simplification, AI will accelerate those goals. Unfortunately, today's healthcare system heavily rewards the former—which, left unmanaged, will lead to more administrative activity, complexity, and spending without commensurate improvements in patient outcomes, access, or affordability.

The early experiences of administrative and clinical AI adoption illustrate this challenge. Applications that fit within today's fee-for-service incentives—particularly in administrative AI—are spreading most rapidly, despite limited evidence of system-wide benefit. In contrast, many clinical AI applications with the greatest potential to improve access, quality, and affordability continue to face significant barriers to adoption, despite growing evidence of their value.


High-value AI measurably improves health outcomes, expands access to care, and lowers spending.

In the near term, adoption of AI-based care should focus on the treatment of health conditions that are access-constrained and have measurable outcomes, such as hypertension or mental health needs. With appropriate safeguards and oversight, these applications of clinical AI stand to deliver the greatest benefits by expanding capacity to deliver care for undertreated populations. Measurable outcomes are important for monitoring the performance of these technologies and for determining appropriate payment. Early use cases should demonstrate that AI can deliver safe and effective care across populations and settings, which is necessary for broader adoption.

Using AI to drive better health and broader system efficiency will require intentionally building an environment that accelerates high-value AI and discourages uses that inflate profits without delivering system-wide benefits. Supporting new AI-enabled models of care delivery will require coordinated

action across multiple market and policy domains, including understanding what works, establishing rules for coverage and use, purchasing effectively, and redesigning care delivery, as illustrated in Exhibit 1.

Exhibit 1
A ROADMAP FOR SCALING HIGH-VALUE AI



Among these levers, payment is foundational because it influences where investment flows, how evidence is generated, which technologies scale, and whether AI is directed toward more activity or better outcomes. By shifting payment over time to models that reward measurable outcomes, we can also redirect administrative activities to those that support patient care and clinical outcomes.

AI presents both a generational opportunity and a generational responsibility. Without deliberate action, AI risks amplifying the inefficiencies and costs of today's healthcare system. Under the right conditions, it can expand clinical capacity to meet a growing demand that cannot be addressed by the current workforce alone, improve outcomes, reduce healthcare spending, modernize the workforce, and simplify administration.

Key Findings

1 AI is fundamentally different from prior health technologies because it can autonomously perform work that previously required human labor and expertise.

2 Under current incentives, administrative AI entrenches inefficient workflows and drives inflationary billing, rather than eliminating waste.

3 AI has the potential to meaningfully improve clinical care, yet adoption remains limited by the current payment, purchasing, and policy environment.

4 Payment reform can accelerate adoption of high-value clinical AI. Outcome-based payment is an important starting point.

About This Report

PHTI convened leaders from health systems, medical societies, health plans, technology companies, academia, investment firms, patient groups, and federal government agencies—including practicing clinicians—for a series of workshops across 2025 and 2026 focused on:

- **Clinical AI:** What evidence, performance monitoring, and regulatory environment will be required for clinical AI to be adopted at scale?
- **Administrative AI:** How can AI be used to simplify administrative processes and reduce spending?
- **Payment for Clinical AI:** How should payment support adoption of AI to improve clinical outcomes while driving down the total cost of care?

Based on discussions during the workshops, as well as supplemental interviews with experts, this report sets out PHTI's recommendations to accelerate the adoption of high-value AI in healthcare that measurably improves health outcomes, expands access to care, and lowers spending.

We conclude that **the healthcare system must intentionally build an enabling environment that directs AI toward measurable improvements in outcomes with clear guidance for coverage and use, while ensuring that AI can be purchased effectively and applied productively** (see Exhibit 1). AI's impact will not be determined by its technical capabilities alone. How it is adopted, what it is rewarded for doing, and the environment in which it operates will ultimately determine the value it delivers.

Without intentional and coordinated action, market forces will continue to steer AI investment and adoption toward activities that reinforce a high-cost, activity-based system—rather than toward innovations that measurably improve care and reduce costs.

Definitions

Clinical AI: AI technologies that assist or perform activities related to the screening, diagnosis, treatment, or management of health conditions.

Administrative AI: AI technologies that assist or perform operational and administrative functions. In this report, administrative AI focuses on transactional activities associated with payment, including prior authorization, coding, and billing.

High-Value AI: AI that measurably improves health outcomes, expands access to care, and lowers spending.

An Opportunity to Transform Healthcare: Why AI Is Unique

AI is fundamentally different from prior health technologies because it can autonomously perform work that previously required human labor and expertise.

AI represents a fundamental shift in how healthcare can be delivered. Most technological innovations to date generally require active human oversight to complete tasks. Today, AI can increasingly perform cognitive, clinical, and administrative work that was historically limited by human time and effort. In clinical care, AI systems can reason through cases and generate diagnostic and treatment recommendations.¹ In transactional activities, AI can document, code, support prior authorization, and conduct patient outreach. AI creates this new clinical capacity in two ways: by enabling the existing workforce to serve patients more effectively and efficiently (assistive AI) and by providing additional clinical services directly to patients (autonomous AI). The same is true for administrative capacity.

AI also makes it possible to move beyond healthcare's traditional reliance on discrete encounters and tasks. Healthcare has largely been organized around visits and the finite time clinicians and staff have available. AI systems can operate continuously, asynchronously, and across large populations, supporting longitudinal care, ongoing patient management, and near-real-time administrative processing. By decoupling output from workforce size and capacity, AI enables the delivery of services at a scale that traditional labor-based models alone cannot reach.

Together, these capabilities introduce a fundamentally new paradigm for healthcare delivery—one in which technology increasingly performs activities at a scale that cannot be achieved by training and hiring more healthcare workers alone. If deployed effectively, AI can expand access to high-quality care, extend workforce capacity, reduce administrative complexity, and improve affordability by enabling services to scale at lower marginal cost.

But this future is not guaranteed. Despite growing evidence that AI can transform care delivery, its adoption has been primarily concentrated in administrative use cases, many of which increase system activity and spending.

Clinical Use Cases

Clinical AI can address persistent healthcare challenges. With expanded AI capacity, care teams can monitor more patients, respond more quickly, and spend more time on the aspects of patient care that require human judgment, personal connection, and individualized decision-making. Early evidence suggests that AI is capable of screening, diagnosis, treatment, and management of health conditions. AI systems can identify patterns that are difficult for clinicians to detect,^{2,3} recognize patient deterioration earlier,⁴ diagnose illness,^{5,6} reduce risk of clinical escalations,⁷ and avoid downstream hospitalizations.⁸

Emerging evidence also suggests that AI systems may improve patient engagement and access by making care available when patients need it and allowing them to interact in ways that are more convenient and comfortable.⁹ Some patients feel that interacting with AI is less stigmatizing than traditional care.^{10, 11}

“The biggest problem with mental health care in America is not access, it's engagement.

People do not like the humans they are meeting. There is data that mental health chatbots can get better engagement than humans.”

As AI capabilities advance, they create the possibility of new care delivery models in which responsibility for activities such as monitoring, triage, diagnosis, treatment, and follow-up is increasingly shared between clinicians and AI systems. Exhibit 2 illustrates how the clinician's role, care delivery, and the potential for scale may evolve as care transitions from traditional clinician-led workflows to AI-assisted care and increasingly autonomous models.

Exhibit 2

ILLUSTRATIVE CLINICAL AI EXAMPLES: EVOLUTION OF AI-BASED CARE DELIVERY MODELS

Level of AI Autonomy	No AI	Assistive AI	Autonomous AI
Clinician Role	Clinician performs all cognitive work	AI performs some cognitive work; clinician reviews and makes decisions	AI performs most cognitive work; clinicians may oversee models and provide input as needed
Care Delivery	Episodic	Continuous monitoring and access; episodic decisions	Continuous monitoring and access
Potential for Scale	Constrained; scale impacted by clinician capacity	Limited; requires clinician review and approval for decisions	High; enables population-level health management

ILLUSTRATIVE CLINICAL EXAMPLES

Hypertension Management	Patient takes blood pressure and brings readings to visits; clinician updates patient's prescription at visit.	Patient takes blood pressure; AI continuously analyzes readings between visits and drafts prescription recommendations. Clinician approves treatment changes.	AI measures and continuously analyzes readings between visits; AI updates patient prescriptions, in most cases, and flags high-risk cases for clinician review.
Mental Health Therapy (Anxiety and Depression)	Patient waits for a scheduled session with a licensed mental health professional who provides therapeutic support at in-person or virtual visits.	Patient sends messages between sessions; AI chatbot drafts responses and symptom summaries. Licensed mental health professional reviews and approves messages.	Purpose-built AI chatbot messages the patient directly between sessions, delivers therapeutic content, and escalates high-risk situations to a licensed mental health professional.
Diabetic Retinopathy Screening	Ophthalmologist performs diabetic retinopathy diagnostic assessment.	AI conducts the initial review of retinal image; ophthalmologist reviews and confirms diagnosis.	AI independently reviews retinal image, provides a diagnostic assessment, and escalates complex or high-risk cases for ophthalmologist review.
Low-Acuity Triage (Urgent Care)	Patient presents in person; clinician performs diagnostic assessment and determines the treatment plan.	AI assesses patient-reported symptoms and drafts a triage decision and treatment plan; clinician reviews and approves.	AI assesses symptoms and delivers a treatment plan to the patient and escalates high-risk cases to a clinician.
Post-Operative Follow-up	Patient attends a scheduled post-operative visit regardless of recovery progress.	AI triages patient-reported symptoms and images and recommends which require follow-ups; clinician reviews and approves.	AI clears uncomplicated recoveries without a visit and escalates high-risk cases to the surgical team.
Inpatient Deterioration Monitoring	Nurse records vital signs and calculates an early warning score; clinician evaluates the patient when the score crosses a threshold.	AI continuously analyzes vitals, labs, and clinical notes and generates a deterioration risk score; care team monitors and determines whether to evaluate the patient.	AI continuously analyzes vitals, labs, and clinical notes and initiates predefined escalation protocols, such as activation of a rapid response team to the patient's bedside.

Administrative Use Cases

Administrative waste is estimated at \$350 billion annually, with \$266 billion attributed to administrative complexity and \$59–\$84 billion to fraud and abuse.¹² AI tools are already being deployed by providers, health plans, and other healthcare organizations to improve efficiency and reduce the cost of administrative processes. For example, providers are using ambient scribing, AI-assisted coding, chart abstraction,

chart summarization, and appeals-drafting tools to automate documentation and administrative workflows. Health plans are applying AI to support prior authorization, claims review, and claims adjudication. Exhibit 3 includes examples of administrative AI currently deployed—most of which automate existing payment and billing transactions.

Exhibit 3

ILLUSTRATIVE ADMINISTRATIVE AI EXAMPLES

Example	Description
AI-Assisted Scribing and Coding Tools	AI transcribes a patient visit, generates a summary, and translates patient conditions into standard codes to support visit documentation and medical billing.
Appeals Drafting	AI reviews a denial letter, pulls the supporting clinical documentation from the chart, and drafts an appeal, citing payer medical necessity criteria.
Chart Abstraction	AI reviews unstructured notes, labs, and imaging reports to pull the specific data elements a clinician or reviewer needs (e.g., prior therapies).
Parsing Complex Plan Policy	AI reviews a payer’s coverage policy and returns the documentation and prior authorization requirements that apply to a given patient and service.
Claims Review	AI reviews submitted claims for coding errors or documentation mismatches before payment, reducing denials and appeals.
Automated Eligibility and Benefits Verification	AI reviews applicant eligibility in real time, eliminating provider and health plan back-and-forth.

Both the clinical and administrative examples illustrate opportunities for AI in healthcare, many of which have the potential to be high-value if deployed under the right conditions. The clinical use cases demonstrate how AI could expand access to care, improve outcomes, and make more productive use of scarce healthcare resources through new models of care delivery. The administrative use cases, by contrast, illustrate the path AI is following today: automating

existing transactions and workflows, some of which are increasing administrative activity and spending.

AI accelerates the processes and functions it is deployed to support. Whether it helps solve healthcare’s persistent challenges or reinforces existing inefficiencies will depend on the incentives that guide where the technology is deployed and scaled.

Accelerating Broken Processes: The Risks of Administrative AI

Under current incentives, administrative AI entrenches inefficient workflows and drives inflationary billing, rather than eliminating waste.

While AI has the potential to reduce complexity and streamline healthcare administration, current applications are often layered on top of broken workflows, lengthening adversarial processes and increasing overall spending. While this may create value for individual organizations (e.g., health plans or hospital systems) by increasing revenue or reducing operational costs, it does not necessarily translate into reduced time and expense for the system as a whole, or even for individual patients. In fact, efficiency and revenue gains at the organization level can result in higher overall healthcare spending by increasing billable activity or intensifying interactions between providers and payers.


Prior Authorization

Prior authorization is the process in which providers get approval from payers before delivering select healthcare services or prescribing medications. The process is estimated to reduce medical and drug spending by up to 5%.^{13, 14} It is also

administratively burdensome, driven by inconsistent requirements across payers and other operational inefficiencies.

The majority (60%) of prior authorization transactions rely on phone- or fax-based workflows.¹⁵ Each prior authorization submission cycle is estimated to cost health plans \$40–\$50 and providers \$20–\$30, as the process often consists of repeated verification, documentation, triage, decision, and appeal cycles.¹⁶

Providers and health plans alike are investing heavily in AI to automate prior authorization workflows to make them faster and less burdensome.¹⁷ For example, on the provider side, AI solutions can draft medical necessity justifications, complete submission forms, and generate appeals; on the health plan side, AI solutions triage requests by complexity, route them to resolution pathways, and provide decision support (Exhibit 4). These tools enable providers to submit more complete requests with less effort and health plans to process higher volumes of submissions at a lower cost per decision.¹⁸



However, the largest sources of administrative burden are not the mechanics of completing prior authorization requests that AI can automate—they are the complexity and variation that underpin the process.

Prior authorization operates within a fragmented system of payer-specific requirements and medical necessity policies. Across four major U.S. insurers, more than 5,000 procedure codes require prior authorization but only 3% of those codes require prior authorization across all four payers.¹⁹ Definitions of medical necessity vary across plans, reviewer discretion can change determinations, and electronic health record workflow limitations make it difficult for providers to access payer-specific requirements at the point of care.

Moreover, costly administrative burden is often concentrated in a relatively small subset of prior authorization requests involving conflicting documentation requirements or disputed interpretations of medical necessity. These cases are difficult to identify at the point of submission, vary across payers, and frequently require iterative review, direct clinician conversations, and appeal cycles.

While AI tools are automating documentation, submission, routing, and review processes, insurers and vendors have not yet reported evidence that AI reduces the average cost per claim after factoring in the cost of the AI solution.²⁰ This reflects two dynamics. AI automation does not address the underlying variation in payer policies that drive the most complex, contentious, and time-intensive prior authorization cases. It also increases the speed of each prior authorization submission, denial, and appeal, creating more back-and-forth

between payers and providers. As a result, AI designed to optimize each side of the transaction makes the overall process more activity-intensive, rather than more efficient.

Policy efforts increasingly focus on standardizing and digitizing prior authorization data and infrastructure. These efforts are an important step toward reducing administrative friction, adding system-level value, and supporting use of AI, but they are unlikely to address the largest sources of administrative burden. Opportunities to build on current efforts include:

- Requiring health plans to publish machine-readable medical policies and escalation pathways to allow providers (and AI tools) to identify requirements before submission.
- Establishing requirements for electronic health record vendor integration of standardized prior authorization application programming interfaces (APIs).
- Requiring clear prior authorization decision rationale, shifting responses from binary approvals and denials to structured next steps (e.g., resubmit the prior authorization with more information or identify an alternative treatment plan for the patient).
- Monitoring denial patterns to identify policy gaps and improve decision responses.

“AI tools are flooding the system.

Bots don't get tired of asking questions, so my review queue keeps growing.”

PRIOR AUTHORIZATION

Market Efforts

Emerging real-time adjudication models, including Optum Real²¹ and the Highmark and Abridge product collaboration,²² aim to complete the full prior authorization at the point of care. This would include integrating clinical documentation, policy matching, and real-time authorization determinations.

Early pilots suggest real-time adjudication models could reduce initial denials and accelerate approvals of prior authorization requests. Self-reported pilot results of Optum Rx's PreCheck Prior Authorization with Cleveland Clinic reported an 88% reduction in appeals and a 68% reduction in denials caused by missing information.²³ Building the evidence base for whether these types of early results can be scaled affordably and broadly into additional clinical settings and contexts is critical.

Currently, the gap between pilots and generalizable solutions is substantial. Real-time adjudication requires reconciling fragmented medical policies across health plans, navigating proprietary medical necessity criteria, managing variation in clinical complexity, and building infrastructure for bidirectional data exchange. Expanding this model from pilots and narrow use cases to a scalable solution across a broad, multistakeholder market will likely require further policy intervention.

PRIOR AUTHORIZATION Policy Efforts

Current policy efforts focus on standardizing and digitizing prior authorization data and infrastructure.

The CMS Interoperability and Prior Authorization Rule (CMS-0057-F) requires select health plans to implement prior authorization APIs by January 1, 2027, enabling providers to determine prior authorization requirements, retrieve documentation templates, and receive decisions electronically in real time. The rule also establishes decision timeframes.²⁴

Building on CMS-0057-F, more than 50 health plans committed to voluntary reforms to simplify and improve the prior authorization process.²⁵ These efforts may reduce friction for straightforward prior authorizations but it is too early to determine impact at scale.

AI Impact on Medical Billing

AI scribing and AI coding tools are now standard across health systems: in a recent survey of large health systems, all reported adopting AI tools for clinical documentation.²⁶ Among clinicians that use AI scribes consistently, emerging evidence suggests that the tools can meaningfully reduce documentation-related cognitive load and burnout.^{27–29}

However, by generating more-complete clinical documentation, these tools increase billing intensity across evaluation and management and diagnostic-related group complexity add-on codes. One multihospital health system found that deploying an AI scribe in its ambulatory and emergency department settings led to a 5% increase in Level 5 encounters and a 7% increase in Level 4 encounters for established patients.³⁰ This translated to an average monthly revenue increase of \$1,004 per provider using an AI scribe.

The full impact of AI-supported documentation and coding tools is not available in claims data given reporting lags. However, health plans describe an immediate increase in coding intensity when AI tools are activated, and major insurers cite aggressive coding as a contributor to higher medical expenditures.^{31, 32} Several plans are deploying algorithms

“ We have a sicker population, and AI is helping us capture more of the details.

The cost will explode unless we bring prices down. It's basic math—there's nothing else that you can do.”

to identify coding outliers, adjust reimbursement, and scrutinize claims.³³ These approaches vary across payers, are not well documented, and their impact on total cost of care is not yet known. AI is accelerating growth in high-complexity billing, which has already strained affordability in recent years.³⁴ The result is a trajectory that payers and patients cannot continue to absorb.

These early experiences with AI in prior authorization and medical billing offer a cautionary lesson about adopting AI: the technology will maximize whatever financial incentives are built into the system. Unfortunately, today's system rewards transactions, documentation, and billing activity, instead of administrative simplification.

Unrealized Potential: The Promise of AI for Clinical Care

AI has the potential to meaningfully improve clinical care, yet adoption remains limited by the current payment, purchasing, and policy environment.


AI enables an evolution in care delivery that expands clinical capacity, improves access to care and health outcomes, and extends high-quality care to underserved populations.

While there is potential for AI to improve care across many conditions, its most immediate applications should be for conditions where access to high-quality care is limited and where outcomes can be measured to monitor performance. Treating hypertension and mental health conditions, such as anxiety and depression, using clinical AI could deliver substantial benefits to patients. For each of these conditions, AI models are increasingly capable of managing and treating patients at low cost with limited human involvement.

In addition, an early focus on conditions with measurable outcomes can provide the foundation for evaluating, improving, and scaling AI-based models of care. Clear clinical endpoints create meaningful goals against which performance can be assessed, enabling developers, providers, payers, and regulators to determine whether AI is delivering value to patients. This accelerates evidence generation and helps build the confidence needed for broader trust, adoption, and, ultimately, payment.

Hypertension is one of the most prevalent and well-understood conditions in the U.S., affecting nearly half of adults.³⁵ Despite the availability of effective and low-cost treatments, only one in four adults with hypertension have adequate blood pressure control.³⁶ Effective hypertension management requires ongoing monitoring, timely medication adjustment, and sustained patient engagement. These activities are difficult to deliver consistently through episodic, visit-based care. AI-based care models can continuously monitor patients, identify those who need intervention, and support treatment decisions between visits. This creates a pathway toward more proactive and scalable population management. Hypertension's well-established and readily measurable endpoints make it possible to clearly assess whether these new models are improving outcomes.

Emerging evidence and commercial developments show that AI-based tools can improve blood pressure control (Exhibit 5). Studies suggest that AI can perform hypertension management tasks, including blood pressure screening and medication management.^{37–39} Companies are also beginning to automate components of hypertension care,⁴⁰ supported by new regulatory and payment pathways for AI-based care.⁴¹



Mental health care represents a similar opportunity for improvement, given that access to services remains severely constrained. More than one in five U.S. adults suffer from mental health symptoms,⁵⁸ yet few receive professional treatment because of workforce shortages, cost barriers, and stigma. In fact, almost half of the U.S. population lives in a Mental Health Professional Shortage Area.⁵⁹


Purpose-built mental health chatbots that meet rigorous standards for safety could help fill these gaps by providing continuous engagement, monitoring symptoms over time, delivering evidence-based interventions, and escalating patients when additional care is needed. While outcomes for anxiety and depression are more complex than for blood

pressure, validated scales that measure symptoms and function provide reliable clinical endpoints that can be used to evaluate the effectiveness of AI-based mental health treatment.

Emerging evidence and commercial developments show that AI-based chatbots can reduce symptoms of depression and anxiety, as well as support sustained patient engagement (Exhibit 6). Clinically specialized AI chatbots can outperform general-purpose large language models (LLMs) on clinician-rated cognitive behavioral therapy competencies.⁶⁰ In practice, however, many users rely on general-purpose LLMs that were not designed for psychotherapy or crisis intervention. These models may reinforce harmful thinking rather than challenge it.⁶¹

Exhibit 6

MENTAL HEALTH CARE: TODAY'S EVIDENCE AND COMMERCIAL FRONTIER



Despite the clinical promise of AI and rapid advances in the underlying technology, significant barriers to adoption remain. Increasingly, however, those barriers are structural rather than technological. AI capabilities are advancing faster than the systems that set the rules and prepare the workforce for healthcare delivery.

Today's healthcare system is built around the assumption that clinical work is performed by humans. Rules define who can deliver care, which services can be reimbursed, how care can be delivered, and who is accountable when something goes wrong. Clinical AI challenges these assumptions, requiring new approaches to authorize its use, integrate it into care delivery, establish payment, and oversee it at scale.

Structural Barriers to AI Adoption

- Limited evidence and immature evaluation frameworks:** Evidence standards, safety thresholds, outcome measures, performance benchmarks, and post-deployment monitoring frameworks are underdeveloped for adaptive and increasingly autonomous AI systems. Regulators, end users, and purchasers lack approaches to determine which technologies work, for whom, and under what conditions. New evaluation approaches are needed to ensure the safety of technologies that evolve over time and perform differently in different environments. Additional testing and evidence generation can help build provider and patient confidence in AI-based care to overcome persistent concerns about safety, reliability, transparency, bias, accountability, and the appropriate role of human oversight.

- **Existing rules for technology coverage and use were not designed for AI:** Existing regulatory, professional, liability, and coverage frameworks were designed for clinician-delivered care and traditional decision-support technology. As AI systems increasingly provide recommendations or perform clinical functions autonomously, unresolved questions remain regarding accountability, liability, oversight, and scope of practice. Today's liability rules generally place responsibility on the supervising clinician or provider organization, creating concerns about accountability for technologies they may not fully understand, control, or develop. At the same time, existing frameworks provide limited clarity regarding how responsibility should be allocated among developers, provider organizations, and clinicians as AI systems assume greater autonomy in care delivery.
- **Payment models were not designed for technologies that independently perform clinical work:** Today's healthcare payment models—fee-for-service, pay-for-performance, and capitation—were not designed for technologies that can independently perform clinical work. As a result, applying today's payment structures to clinical AI creates a structural misalignment. These approaches do not readily accommodate technologies that automate clinical work, assist clinician decision-making, or operate with increasing autonomy. In fact, applying fee-for-service payment to AI risks increasing spending. Establishing sustainable payment for autonomous clinical AI is necessary to realize system savings, drive further technological innovation, encourage adoption by providers and patients, and justify the implementation effort required for new solutions.

- **Clinical and administrative workflows do not readily integrate autonomous or AI-assisted processes:** Healthcare organizations continue to rely on care delivery models, administrative processes, workforce structures, and governance approaches that were developed for a pre-AI environment. Integrating AI into care delivery will require modernized workforce and organizational capabilities.⁷⁴ Leadership and staff roles, workflows, education and training, and oversight structures need to be redesigned to work alongside AI systems. This shift will also require changes to education, training, and workforce development. Delivery models will need to restructure the allocation of tasks to maximize the benefits of AI, while deploying human effort to the highest-value tasks.

While addressing each of these barriers is necessary to drive AI adoption, payment is a particularly important first step. It determines whether organizations can justify the investments needed to adopt and scale AI in the first place. Without a reliable payment pathway, providers, investors, researchers, and developers have limited incentives to absorb the costs of activities like evidence generation, workflow redesign, governance, workforce training, and performance monitoring. As a result, payment is one of the most important determinants of whether clinical AI moves from technical promise to widespread adoption.


Setting the Right Incentives: How Payment Models Drive AI Adoption

Payment reform can accelerate adoption of high-value clinical AI, with outcome-based payment as an important starting point.

Payment is the primary mechanism through which the healthcare system communicates what it values. It determines where investment flows and which innovations scale. Organizations invest in technologies when payment and incentives support adoption.

The Limitations of Today's Payment Models

Fee-for-service remains the dominant payment model today. It was designed to pay for human labor, time, and discrete activities—not for technology (Exhibit 7). It assumes that delivering more care requires more clinician effort, and, thus, more payment.



Note: Health Care Payment Learning & Action Network (LAN), [2020–2021 Alternative Payment Model \(APM\) Measurement Report](#). *60% includes FFS with a link to quality and value; 85%–90% of services rely on FFS including APMs built on FFS architecture.

Exhibit 8

STRUCTURAL MISMATCH BETWEEN FEE-FOR-SERVICE PAYMENT AND CLINICAL AI

Today's Fee-for-Service Payment Inputs

Structural Mismatch with Clinical AI

LABOR-BASED PAYMENT



AI decouples clinical output from clinician effort and time.

ENCOUNTER-BASED PAYMENT



AI enables continuous, patient-directed longitudinal care.

UTILIZATION-BASED ECONOMICS



AI marginal cost approaches token consumption costs at scale.

These assumptions break down in the context of clinical AI. Effective clinical AI generates more clinical output without a proportional increase in clinician labor; delivers clinical value without consuming clinician time; and enables continuous, patient-directed care rather than discrete encounters (Exhibit 8). At scale, because marginal costs resemble those of software rather than labor, clinical activity can expand without a commensurate increase in spending.

The mismatch grows more pronounced as AI autonomy increases (Exhibit 9). Fee-for-service could partially accommodate assistive clinical AI because a clinician remains heavily involved in care delivery. Existing supervisory billing codes could be modified to allow clinicians to bill for encounters involving oversight of AI-generated clinical recommendations. AI used to interpret mammograms or chest x-rays, for example, can be reimbursed today under existing Category 1 CPT codes.

As technology performs a growing share of clinical work autonomously, there are fewer established payment options and unclear mechanisms for determining how services should be valued and paid for. It also raises questions about whether and how technology developers or operators could be paid directly.

Today, providers deploying clinical AI primarily bill through existing payment pathways and absorb the costs of these tools. This approach may support limited adoption but is unlikely to scale without a clear payment mechanism.

There may be focused use cases where fee-for-service payment models for autonomous AI could be high-value, despite the inflationary risk. As an example, fee-for-service payment may be suited for autonomous diagnostics, where there is a necessary task to be completed without a clear link to outcomes. However, in most cases, fee-for-service creates the wrong incentives for clinical AI. Because AI can scale at low marginal cost, paying for each service or interaction can increase volume of billable activity without necessarily better outcomes.

Exhibit 9

STRUCTURAL FEE-FOR-SERVICE PAYMENT MISMATCH INCREASES WITH LEVEL OF AUTONOMY

Example: Hypertension Care Model

Level of AI Autonomy	No AI	Assistive AI	Autonomous AI
	Patient manually tracks blood pressure and a clinician reviews the data to create a titration plan at each visit.	AI continuously monitors blood pressure and drafts titration recommendations, which a clinician must approve before action.	AI continuously analyzes blood pressure and delivers medication titrations to the patient with minimal clinician involvement.
Labor-Based Payment	Clinician work drives care; payment maps to time spent with and complexity of patient	AI may compress clinician effort and time	Minimal clinician work involved; oversight only
Encounter-Based Payment	Care is mostly episodic, visit-focused	AI performs activity between encounters	AI monitoring and titration replace encounters
Utilization-Based Payment	Volume limited by clinician capacity	AI may enable increased volume of activity per clinician and increase costs	AI-created potential for cost inflation with near-zero marginal costs and unlimited potential for scale
Overall Fit with FFS Payment Model	 GOOD	 PARTIAL	 POOR

Performance-based and capitated models are better suited to AI because they shift incentives from volume toward value and outcomes. But as designed today, these payment models do not create sufficient incentives to justify the workflow redesign, operational changes, or upfront investment required to support adoption of clinical AI.

Adoption will be driven less by technical capability than by what the system pays for. Healthcare faces a choice: maintain a cycle in which limited payment constrains deployment—slowing evidence generation and trust—or develop payment models that accelerate both evidence development and adoption.

The Case for Outcome-Based Payments for Clinical AI

Healthcare delivery and administration are largely organized around billable services. Clinical workflows generate reimbursable encounters, tests, procedures, and prescriptions, while administrative workflows support the documentation, coding,

prior authorization, claims processing, and reconciliation required to receive payment. AI has advanced most rapidly where it reinforces existing structures. AI scribing, coding tools, claims-processing applications, and prior authorization solutions reduce manual effort, accelerate transactions, strengthen documentation, and improve revenue cycle performance, offering organizations a clear and immediate return on investment.

Many promising clinical AI applications create value differently. Rather than enabling additional billable activities, the technology may reduce clinician time, automate care delivery typically conducted by a clinician, or enable more continuous and scalable care. Although these technologies could improve outcomes, expand access, and lower total costs, their value does not translate into dependable revenue under fee-for-service payment built around clinician time and discrete encounters. As a result, the market rewards AI that optimizes transactions more readily than AI that transforms care delivery.

“Implementing new technologies requires significant change management.”

Even in ACOs, where the incentives are aligned to support technology adoption, we found clinician resistance to changing the way they work.”

Even in capitated and full-risk environments, where incentives are aligned toward reducing total cost of care, adoption of digital health technologies has been slow. A national survey of 276 organizations with Medicare or Medicaid accountable care organization contracts found that only one-third adopted digital health technology for a high-priority, pervasive chronic condition, and those that did used it to complement rather than replace traditional care.^{75, 76} Without deliberate payment and workflow redesign, clinical AI is likely to face many of the same barriers to adoption.

Outcome-based payment represents an important starting point for shifting investment toward high-value clinical AI. When payment is tied to measurable results, the incentive is to improve outcomes and reduce spending. This creates a business case for overcoming the structural barriers to clinical AI adoption, including evidence generation, workflow redesign, and organizational change. Organizations are more likely to invest in this work when there is a clear path to compensation for outcomes and cost savings generated through AI-based care. Without such pathways, healthcare may never be able to determine whether AI will allow us to deliver better outcomes at lower costs.

Importantly, AI-based care models allow outcomes to be measured beyond an individual encounter level. These models create value through continuous monitoring, treatment optimization, and longitudinal disease management. The outcomes that AI systems help produce may be best assessed across episodes of care and throughout a patient’s longitudinal care journey, rather than through a single visit or intervention.

The CMS Innovation Center’s ACCESS (Advancing Chronic Care with Effective, Scalable Solutions) outcome-based payment model does not require the use of AI, but its structure accommodates and encourages technology in ways that traditional payment models do not. Because ACCESS rates are relatively low, there is a strong incentive for entrants to leverage highly automated or AI-based tools to support patients at a lower cost. ACCESS also requires participants to report outcomes through standardized APIs. This creates a real-world, national mechanism for AI-based care models to compete and be evaluated. Ultimately, healthcare must establish pathways to evaluate the ability of AI systems to deliver improvements in outcomes and efficiency. Payment models tied to real-world value provide an important starting point to accelerate this shift.

However, no single payment model will support AI-enabled care across all use cases and clinical settings. Outcome-based payment, for example, is not well suited to settings where outcomes are difficult to measure or attribute to a specific intervention, such as primary care. For some use cases, fee-for-service payment with a built-in price reduction mechanism may be the most appropriate option. Further work is needed to design payment models for a range of AI-enabled care applications.

How Payment Reform Could Redirect AI Toward Administrative Simplification

Payment reform can change today’s incentives and redirect AI toward administrative simplification. Under fee-for-service payment, the system must verify that each service happened and that it was medically necessary. This creates administrative processes including prior authorization, time-based tracking and reporting, encounter-level documentation, coding review, claims adjudication, and appeals.

Under outcome-based payment models, the administrative activity associated with each billable service becomes less central. Instead, administrative functions shift toward processes that support care delivery and accountability, such as outcomes measurement, quality reporting, population attribution, and risk adjustment with the goal of improving patient care and performance.

The Path Forward: Proactive Development of High-Value AI

AI presents a generational opportunity to improve health outcomes, expand access to care, and reduce healthcare spending. It also creates a responsibility to ensure that policy and market incentives direct investment and create the right ecosystem for technologies that produce meaningful, system-level value—not simply more transactions and spending.

To achieve this vision for high-value AI in healthcare, policymakers and market leaders must advance a coordinated approach to build conditions that ensure clinical and administrative AI will add value to society (Exhibit 10). These efforts must include

identifying which AI-based care models are effective, establishing rules for coverage and use, effective purchasing, and redesigning care delivery and administrative workflows.

Exhibit 10

A ROADMAP FOR SCALING HIGH-VALUE AI



The roadmap is organized around four domains, each of which must advance to build an enabling environment for high-value AI adoption:

DOMAIN

1

Understand what works: Invest in an evidence ecosystem to understand which clinical AI tools work, for whom, and under what conditions. This will require developing evidence standards, defining outcomes measures, and building benchmarking frameworks that allow purchasers, researchers, and end users to evaluate performance across tools, populations, and settings over time.

- **Generate real-world evidence:** AI may perform differently in different populations, settings, and workflows, or when used by clinicians or staff with different levels of training and resources. While studies in simulated settings suggest AI may outperform physicians at specific clinical tasks,⁷⁷ evidence that AI can manage care and improve outcomes in real-world settings remains limited.^{78, 79} Initial testing and ongoing real-world evidence generation are essential to determine whether a tool improves outcomes in practice. Federal and state demonstrations, regulatory sandboxes, and commercial pilots should allow promising technologies to be tested with clear safeguards. These initiatives should have defined evaluative endpoints against an established

comparator. These programs should also identify the evidence required to advance beyond the pilot and the conditions under which testing will be modified or stopped.

- **Assess outcomes and monitor ongoing performance:** Providers, payers, associations, and healthcare quality organizations should identify specific outcome measures and define the level of improvement needed for AI technologies to succeed within a given clinical context. These measures should focus on meaningful clinical, operational, patient experience, and spending outcomes, rather than on process measures. Evaluation frameworks should enable transparent and comparable assessment of AI technologies while allowing benchmarks, comparators, and success thresholds to reflect real-world needs. In most cases, AI should be evaluated relative to current standards of care and existing access to services, rather than to idealized care models that may not be attainable in practice. Target metrics can help focus efforts and provide innovators and investors with greater clarity regarding the outcomes that matter most for adoption and scale. Clear assessment leads to improved evidence generation and more confidence in solution performance by the entities that purchase and use the technology, including clinicians and patients. Further, continued monitoring (e.g., postmarket surveillance

EXAMPLES OF PROGRESS

Generate Real-World Evidence

State Regulatory Sandboxes: The Utah Office of Artificial Intelligence Policy's (OAIP) AI regulatory sandbox approved a pilot that allows an AI system to provide routine, low-risk prescription renewals without any human involvement for patients with chronic conditions.⁸⁰ OAIP subsequently published evidentiary expectations for healthcare AI sandbox participants, including benchmarking performance against independent clinician assessments and establishing criteria for adjusting oversight and pilot continuation.⁸¹ Other states are exploring or are in the process of developing sandbox initiatives, including Delaware⁸² and Connecticut.⁸³

FDA: The FDA is testing new approaches to the oversight of digital health technologies through the voluntary Technology-Enabled Meaningful Patient Outcomes (TEMPO) pilot, launched in coordination with CMS's ACCESS model.⁸⁴ TEMPO creates a pathway for select digital health devices supporting care under ACCESS to reach patients through a more flexible approach to FDA oversight of digital health devices, while participating manufacturers collect and share real-world data on their performance. The pilot focuses on technologies addressing cardio-kidney-metabolic, musculoskeletal, and behavioral health conditions. By generating evidence on how these technologies perform in real-world settings, TEMPO could help inform future approaches to device oversight that support technology innovation while maintaining appropriate safeguards.

Advanced Research Projects Agency for Health (ARPA-H): ARPA-H's Agentic AI-Enabled Cardiovascular Care Transformation (ADVOCATE) program is testing the frontier of clinical autonomy by funding the development of an FDA-authorized cardiovascular agent capable of taking actions alongside supervisory AI designed to monitor its performance. ARPA-H committed \$62.7 million to the program over four years.⁸⁵

EXAMPLES OF PROGRESS

Assess Outcomes and Monitor Ongoing Performance

FDA: In August 2026, the FDA released a discussion paper seeking feedback on generative AI-enabled medical devices, specifically focused on risk assessment, premarket evaluation, postmarket monitoring, and other relevant areas.⁸⁶ The paper highlights the FDA's recognition that traditional evaluation approaches may be insufficient for generative AI technologies and explores new frameworks for assessing safety, effectiveness, autonomy, and real-world performance across the total product lifecycle. The agency proposes a competency-based evaluation model, enhanced postmarket monitoring, and risk-based oversight.

for FDA-regulated products) is necessary to detect performance degradation, unintended consequences, disparities, and emerging risks as clinical practices, patient populations, data inputs, evidence, and the tools evolve.

- **Build patient, provider, and purchaser confidence:** Confidence in AI-based care relies on clear, transparent evidence of what works, for whom, and under what conditions. Early use cases should demonstrate that AI can deliver safe and effective care across populations and settings, which is necessary for broader adoption. Findings from real-world evidence generation, outcomes assessment, and performance monitoring should be reported transparently to support informed adoption and purchasing decisions.

DOMAIN 2 **Set rules for coverage and use:** AI technologies that demonstrate value should have a pathway from testing to coverage. Policies should support responsible deployment while additional evidence is generated.

At the same time, policymakers, professional societies, regulators, and payers will need to establish clear rules governing the appropriate use of these technologies.

- **Update licensure and scope of practice:** Existing rules were designed around care delivered directly by humans and do not fully account for technologies that provide recommendations, perform discrete clinical tasks, or operate with increasing autonomy. Policymakers and professional bodies must establish clear standards for human oversight,

model supervision, escalation, and accountability, with requirements proportionate to a technology's autonomy and risk. Regulatory frameworks will also need to address when and how AI can be authorized to perform clinical functions traditionally reserved for licensed professionals. Lessons from AI regulatory sandboxes, pilot initiatives, and other real-world deployment can help inform future updates to licensure requirements, scope-of-practice rules, and professional standards for AI-based care delivery.

- **Clarify liability and insurance:** Uncertainty regarding liability remains a barrier to the adoption of clinical AI. For example, stakeholders need clarity regarding responsibility when AI makes a mistake or when a clinician ignores an AI-generated recommendation. Today's malpractice frameworks were not designed for care models in which clinical decisions may be shared among clinicians and AI technology. How liability should be allocated in these settings remains unsettled.⁸⁸ New frameworks need to address questions related to oversight and supervision, documentation of decision-making, malpractice, and allocation of responsibility between the developer, clinicians, or provider organization, depending on the level of autonomy.
- **Define coverage across payers:** Policies should provide a clear coverage pathway for high-value AI technologies and ideally be aligned across payers. Payment demonstrations and coverage policy should work together to support responsible deployment. Coverage decisions should include mechanisms to reassess performance over time and allow policymakers and payers to adjust payment levels, impose additional safeguards, limit use, or withdraw support when value declines or risks emerge. Coverage and payment should become more durable only when technologies demonstrate performance in real-world settings.

EXAMPLES OF PROGRESS

Update Licensure and Scope of Practice

Federation of State Medical Boards: Professional organizations are beginning to define how governance, liability, and accountability should evolve as AI assumes more clinical responsibility. In 2026, the Federation of State Medical Boards established a workgroup to develop recommendations for state medical boards, focusing on AI performing clinical functions with limited or no physician supervision.⁸⁷

DOMAIN

3

Purchase effectively: Even when clinical AI demonstrates value, adoption depends on reliable reimbursement and incentives to invest in the technology and associated operational and change management costs. Different AI applications may require different payment approaches depending on the clinical use case, degree of autonomy, and value being created. Outcome-based payment models are an important place to start because they align reimbursement with demonstrated improvements in patient outcomes. Purchasers, payers, and policymakers should evaluate vendors against established outcome measures and develop payment approaches that reward performance. Payment for AI is explored in greater detail in the following section.

DOMAIN

4

Redesign care delivery: Clinical AI will not achieve its full potential if it is layered onto existing healthcare processes. Organizations will need to rethink care delivery models, administrative workflows, and the workforce to take advantage of AI's capabilities while maintaining appropriate oversight, accountability, and human judgment.

- **Incorporate clinical AI into care models:** Organizations and developers will need to rethink how care is organized and delivered as AI enables new approaches to continuous monitoring, decision support, patient engagement, administrative functions, and autonomous clinical care. This will include shifting from episodic encounters to continuous management and expanding population-based care models.
- **Evolve clinical workforce roles and responsibilities:** Healthcare organizations will need to redesign staffing models and train their clinical workforce in new skills. Integrating AI into care delivery requires clinicians to work effectively alongside AI and take on new responsibilities, including supervising AI, validating AI outputs, coordinating multiple models, and overseeing care across patient populations. Organizations will need to determine which functions should remain human, which can be delegated to AI, and how roles and governance should evolve as that division of labor changes. They will also need to address the risks of deskilling through new models of clinical education and training, ensuring that critical human strengths are preserved.

- **Eliminate adverse administrative processes:** AI creates opportunities to simplify workflows and reduce spending on administrative waste. Organizations can use AI to reduce manual work, improve efficiency, and lessen administrative burden. Efforts should prioritize eliminating unnecessary complexity and reducing system-wide waste rather than accelerating the level of activity through automation. This will require coordination among providers, payers, vendors, and policymakers to ensure automation reduces rather than exacerbates friction.

EXAMPLES OF PROGRESS

Advance Outcome-Based Payment

CMS: CMS's ACCESS model is the most concrete step to date toward a payment model that accommodates AI. The 10-year, outcome-based payment model launched in July 2026 with more than 160 participating organizations, focusing on the management of chronic cardiometabolic, musculoskeletal, and behavioral health conditions.⁸⁹ In spring 2027, CMS will expand ACCESS with new tracks for heart failure, chronic obstructive pulmonary disease, substance use disorder, and tobacco cessation.⁹⁰

EXAMPLES OF PROGRESS

Streamline Administrative Processes

CMS and ONC: CMS and ONC (Office of the National Coordinator for Health Information Technology) are advancing digital infrastructure to support AI-based care. Efforts to expand Fast Healthcare Interoperability Resources (FHIR) or FHIR-based interoperability, standardized data exchange, and API-driven performance reporting are creating foundational capabilities needed to support outcomes measurement, evidence generation, and technology-enabled care models at scale.^{91, 92}

Paying for AI in Healthcare: Three Principles to Drive Value

Payment policy will shape which AI technologies are developed, adopted, and scaled. Developing new payment models can steer AI investment and adoption toward technologies that improve outcomes, simplify administrative activity, and reduce healthcare spending. To achieve this, three core principles should guide the design of future payment models for AI-enabled care.

1 Payment must produce system-level value by lowering spending while maintaining or improving clinical outcomes. Payment for clinical AI should be justified only when AI demonstrably improves upon current care, such as by producing better clinical outcomes, lowering spending, expanding access to effective services, or achieving a meaningful combination of these goals.

Value should be measured against what would have occurred without the technology—not against an idealized standard. Evaluations should account for the full impact on healthcare spending, including whether savings in one area are offset by additional testing, referrals, treatment, or administrative activity elsewhere. They should also examine population-level effects, recognizing that a low-cost intervention applied across millions of patients can generate substantial aggregate spending even when the cost per use is small.

Clinical AI should make the healthcare system more productive: delivering better or more accessible care with fewer total resources and lowering premiums and out-of-pocket costs for patients and purchasers.

2 Payment should be tied to results and reward appropriate utilization. AI fundamentally changes the relationship between clinical activity and value. Once developed, an AI system may be able to perform or support millions of clinical interactions at a very low marginal cost. Paying separately for every use can therefore create strong incentives for unnecessary utilization and allow spending to grow far faster than clinical value.

Payment for clinical AI should instead reflect realized clinical or economic outcomes rather than time, effort, or resource inputs.

Payment models should also incorporate explicit safeguards against overuse, underperformance, and adverse patient selection. These could include performance-based withheld payments or price adjustments as utilization expands to prevent overutilization and unconstrained spending. The objective is not to simply pay for the use of AI, but to reward appropriate deployment of the technology that generates measurable value.

3 Payment must reward early innovation and adoption, then evolve with evidence, performance, and the changing clinician role. Payment systems for AI should be more dynamic than those designed for traditional care delivery. Rates should start high enough to stimulate early innovation and adoption, then adjust as utilization scales, marginal costs decline, and real-world evidence accumulates. Rates and payment should be reviewed regularly rather than becoming permanent based on early assumptions. Payment could decline as technologies scale, increase when evidence demonstrates greater value than initially expected, or end when a service fails to deliver meaningful results. While reducing payment is operationally and politically complex, once technologies have matured and evidence demonstrates that AI enables care to be delivered with fewer resources, payment should reflect that new cost structure. Medicare's Physician Fee Schedule provides a precedent.

Payment must also account for and fairly compensate for a rapidly evolving clinician role. As AI becomes more capable of autonomous tasks, clinicians may spend less time directly producing individual services and more time supervising AI systems, validating outputs, managing exceptions, coordinating care, and remaining accountable for the quality of care delivered across a population. Payment policy should incentivize the shift from individual encounters toward population-level management.

No single payment model will support AI adoption across all clinical settings and use cases. Different applications of AI generate different system value, require different levels of clinician oversight, and impact patient outcomes over different time horizons. Instead, new payment models for clinical AI must be deliberately designed for specific contexts.

In the future, PHTI will develop more specific recommendations for designing payment models that apply these three principles across different settings and applications of AI. For instance, while outcome-based payment models are well suited to chronic condition management where there is consensus on established clinical endpoints, AI-enabled care delivered in urgent or primary care settings may require a different payment approach. PHTI will explore a range of use cases that may not lend themselves to outcome-based payments, as well as questions on how to set pricing and how to reimburse providers for oversight of AI systems.

The healthcare system must intentionally build an enabling environment that directs AI toward measurable improvements in outcomes, access, affordability, and administrative simplification.

AI's impact will not be determined by its technical capabilities alone. How it is adopted, what it is rewarded for doing, and the environment in which it operates will ultimately determine the value it delivers.

Looking Ahead

We are at a crossroads: AI could either help solve healthcare's longstanding challenges or accelerate the forces that created them. The work ahead is to build the conditions that will enable high-value AI. This is particularly urgent as clinician shortages deepen and the U.S. population ages. The choices we make today will determine whether AI becomes another layer of cost and complexity or a catalyst for a more accessible, affordable, and outcomes-oriented healthcare system. Realizing this potential will require policymakers, industry leaders, providers, and researchers to modernize payment models, strengthen evidence and accountability frameworks, purchase effectively, and redesign care delivery. Ultimately, lower spending should translate to savings for patients, employers, and taxpayers and lead to a higher-value healthcare system for all.

References

- ¹ Ezekiel J. Emanuel, Abe Baker-Butler, Neal Khosla, and Vinod Khosla, “Will Autonomous AI Exceed AI-Aided Physicians as the Best Medical Care?” *JAMA* 336, no. 11 (2026): 915–918.
<https://doi.org/10.1001/jama.2026.15380>
- ² Edward Korot, Nikolas Pontikos, Xiaoxuan Liu, et al., “Predicting Sex from Retinal Fundus Photographs Using Automated Deep Learning,” *Scientific Reports* 11 (2021): 10286.
<https://doi.org/10.1038/s41598-021-89743-x>
- ³ Xiaoxi Yao, David R. Rushlow, Jonathan W. Inselman, et al., “Artificial Intelligence–Enabled Electrocardiograms for Identification of Patients with Low Ejection Fraction: A Pragmatic, Randomized Clinical Trial,” *Nature Medicine* 27 (2021): 815–819.
<https://doi.org/10.1038/s41591-021-01335-4>
- ⁴ Lo Pang-Yun Ting, Hong-Pei Chen, An-Shan Liu, et al., “Early Detection of Patient Deterioration from Real-Time Wearable Monitoring System,” preprint, *arXiv* (2025): 2505.01305.
<https://doi.org/10.48550/arXiv.2505.01305>
- ⁵ Gina Kolata, “Doctors Thought It Was Asthma. A.I. Flagged a Serious Heart Problem,” *The New York Times*, June 22, 2026.
<https://www.nytimes.com/2026/06/22/health/artificial-intelligence-heart-damage.html>
- ⁶ Jennifer I. Lim and Theodore Leng, “Artificial Intelligence for Diabetic Retinopathy,” *JAMA Ophthalmology*, published online September 3, 2026.
<https://doi.org/10.1001/jamaophthalmol.2026.3479>
- ⁷ Robert J. Gallo, Lisa Shieh, Margaret Smith, et al., “Effectiveness of an Artificial Intelligence–Enabled Intervention for Detecting Clinical Deterioration,” *JAMA Internal Medicine* 184, no. 5 (2024): 557–562.
<https://doi.org/10.1001/jamainternmed.2024.0084>
- ⁸ Ann C. Raldow, Naveen Raja, Chad W. Villaflores, et al., “Proactive Care Management of AI-Identified At-Risk Patients Decreases Preventable Admissions,” *American Journal of Managed Care* 30, no. 11 (November 2024): 548–54.
<https://doi.org/10.37765/ajmc.2024.89625>
- ⁹ Michael V. Heinz, Daniel M. Mackin, Brianna M. Trudeau, et al., “Randomized Trial of a Generative AI Chatbot for Mental Health Treatment,” *NEJM AI* 2, no. 4 (2025).
<https://doi.org/10.1056/Aloa2400802>
- ¹⁰ Clare E. Palmer, Emily Marshall, Edward Millgate, et al., “Combining Artificial Intelligence and Human Support in Mental Health: Digital Intervention with Comparable Effectiveness to Human-Delivered Care,” *Journal of Medical Internet Research* 27 (2025): e69351.
<https://doi.org/10.2196/69351>
- ¹¹ Johanna Habicht, Sruthi Viswanathan, Ben Carrington, et al., “Closing the Accessibility Gap to Mental Health Treatment with a Personalized Self-Referral Chatbot,” *Nature Medicine* 30, no. 2 (2024): 595–602.
<https://doi.org/10.1038/s41591-023-02766-x>
- ¹² William H. Shrank, Teresa L. Rogstad, and Natasha Parekh, “Waste in the US Health Care System: Estimated Costs and Potential for Savings,” *JAMA* 322, no. 15 (2019): 1501–1509.
<https://doi.org/10.1001/jama.2019.13978>
- ¹³ Fritz Busch and Stacey Muller, “Potential Impacts on Commercial Costs and Premiums Related to the Elimination of Prior Authorization Requirements,” Milliman, March 30, 2023.
https://media.milliman.com/v1/media/edge/images/millimaninc5660-milliman6442-prod27d5-0001/media/Milliman/PDFs/2023-Articles/8-18-23_BCBSA-Prior-Authorization-Impact.pdf
- ¹⁴ Zarek C. Brot-Goldberg, Samantha Burn, Timothy Layton, and Boris Vabson, “Rationing Medicine Through Bureaucracy: Authorization Restrictions in Medicare,” NBER Working Paper, no. 30878, National Bureau of Economic Research, January 2023.
<https://doi.org/10.3386/w30878>
- ¹⁵ CAQH, “The 2025 CAQH Index Report,” DataSpring, 2026.
<https://www.dataspring.com/advisory-services/index-report>
- ¹⁶ Nikhil R. Sahni, Brooke Istvan, Celia Stafford, and David Cutler, “Perceptions of Prior Authorization Burden and Solutions,” *Health Affairs Scholar* 2, no. 9 (2024): qxae096.
<https://doi.org/10.1093/haschl/qxae096>
- ¹⁷ Megan Zweig, Mihir Somaiya, and Tiffany Marie Ramos, “H1 2025 Market Overview: Proof in the Pudding,” Rock Health, July 7, 2025.
<https://rockhealth.com/insights/h1-2025-market-overview-proof-in-the-pudding/>
- ¹⁸ Michelle M. Mello, Artem A. Trotsyuk, Abdoul Jalil Djiberou Mahamadou, and Danton Char, “The AI Arms Race in Health Insurance Utilization Review: Promises of Efficiency and Risks of Supercharged Flaws,” *Health Affairs* 45, no. 1 (2026): 6–13.
<https://doi.org/10.1377/hlthaff.2025.00897>
- ¹⁹ Aya Zaari Jabri, Jacob Asher, Jacqueline Sandling, et al., “Variation and Standardization in Prior Authorization Requirements,” preprint, *medRxiv*, April 13, 2025.
<https://doi.org/10.1101/2025.02.07.25321895>
- ²⁰ Mello et al., “The AI Arms Race,” 6–13.

References

- 21 Optum, “Optum Is Advancing AI-Powered Digital Prior Authorization to Reshape Care Delivery,” press release, February 4, 2026.
<https://www.optum.com/en/about-us/news/page.hub5.ai-powered-digital-prior-authorization.html>
- 22 Abridge, “Highmark Health, Abridge Announce Collaboration to Deploy AI Technologies Across an Entire Payer-Provider Ecosystem,” press release, August 12, 2025.
<https://www.abridge.com/press-release/highmark-health-ahn-abridge-prior-authorization>
- 23 UnitedHealth Group, “Optum Rx Automates Prior Authorization Process for Prescription Drugs to Improve the Patient and Provider Experience,” October 22, 2024.
<https://www.unitedhealthgroup.com/newsroom/posts/2024/2024-10-22-optumrx-prior-authorization-process-to-improve.html>
- 24 Centers for Medicare & Medicaid Services (CMS), “Interoperability and Prior Authorization Final Rule (CMS-0057F),” *Federal Register* 89, no. 27 (February 8, 2024): 8758.
<https://www.govinfo.gov/content/pkg/FR-2024-02-08/pdf/2024-00895.pdf>
- 25 America’s Health Insurance Plans, “Health Plans Take Action to Simplify Prior Authorization,” press release, June 23, 2025.
<https://www.ahip.org/news/press-releases/health-plans-take-action-to-simplify-prior-authorization>
- 26 Eric G. Poon, Christy Harris Lemak, Juan C. Rojas, et al., “Adoption of Artificial Intelligence in Healthcare: Survey of Health System Priorities, Successes, and Challenges,” *Journal of the American Medical Informatics Association* 32, no. 7 (2025): 1093–1100.
<https://doi.org/10.1093/jamia/ocaf065>
- 27 Kristine D. Olson, Daniella Meeker, Matt Troup, et al., “Use of Ambient AI Scribes to Reduce Administrative Burden and Professional Burnout,” *JAMA Network Open* 8, no. 10 (2025): e2534976.
<https://doi.org/10.1001/jamanetworkopen.2025.34976>
- 28 Anand Chowdhury, Michele Casey, Jonathan Wilson, et al., “Comparing Ambient Scribes: A Randomized Crossover Clinical Trial Addressing Ambient Scribe Technologies’ Impact on Physician Burnout,” *Journal of the American Medical Informatics Association* 33, no. 5 (2026): 990–999.
<https://doi.org/10.1093/jamia/ocag018>
- 29 Peterson Health Technology Institute (PHTI), “Adoption of Artificial Intelligence in Healthcare Delivery Systems: Early Applications and Impacts,” March 25, 2025.
<https://phti.org/adoption-ai-healthcare-delivery/>
- 30 Tyson Blauer, Mac Boyter, and Sidnee Wood, “Suki ROI Validations 2026: Cross-Organizational Results from Suki’s Clinical Intelligence Platform,” *KLAS Research*, January 2026.
<https://go.suki.ai/Suki-KLAS-ROI-Validations-2026>
- 31 Katie Palmer, “Insurers Say They’ll Deploy More AI to Combat ‘Aggressive’ Coding by Hospitals,” *STAT News*, October 31, 2025.
<https://www.statnews.com/2025/10/31/health-insurer-ai-fights-hospital-ai-coding-billing/>
- 32 BlueCross BlueShield, “Analysis Examines AI in Hospital Billing As Spike in Complex Patients Adds Nearly \$1 Billion in Extra Costs,” press release, September 24, 2026.
<https://www.bcbs.com/about-us/association-news/bcbsa-analysis-ai-coding-tools-affects-healthcare-costs>
- 33 Elizabeth Casolo, “How 3 Insurers Are Trying to Tighten E/M Billing Oversight,” *Becker’s Payer Issues*, March 4, 2026.
<https://www.beckerspayer.com/payer/how-3-insurers-are-trying-to-tighten-e-m-billing-oversight/>
- 34 Lynne Cotter, Emma Wager, Hattie Xu, et al., “Eight Trends Shaping 2026 Healthcare Costs,” *Peterson-KFF Health System Tracker*, March 17, 2026.
<https://www.healthsystemtracker.org/chart-collection/eight-trends-shaping-2026-healthcare-costs/>
- 35 Centers for Disease Control and Prevention (CDC), “Estimated Hypertension Prevalence, Treatment, and Control Among U.S. Adults,” Million Hearts, last reviewed May 12, 2023.
<https://millionhearts.hhs.gov/data-reports/hypertension-prevalence.html>
- 36 CDC, “Estimated Hypertension Prevalence, Treatment, and Control.”
- 37 David J. Collier, Mike Taylor, Thomas Godec, et al., “Personalized Antihypertensive Treatment Optimization with Smartphone-Enabled Remote Precision Dosing of Amlodipine during the COVID-19 Pandemic (PERSONAL-CovidBP Trial),” *Journal of the American Heart Association* 13, no. 4 (2024): e030749.
<https://doi.org/10.1161/JAHA.123.030749>
- 38 Jiayi Yi, Lili Wang, Jiali Song, et al., “Development of a Machine Learning-Based Model for Predicting Individual Responses to Antihypertensive Treatments,” *Nutrition, Metabolism and Cardiovascular Diseases* 34, no. 7 (2024): 1660–1669.
<https://doi.org/10.1016/j.numecd.2024.02.014>
- 39 Yijun Wang, Wuping Tan, Siyi Chen, et al., “Large Language Model Agent for Managing Patients with Suspected Hypertension,” *Hypertension* 83, no. 1 (2026): 212–224.
<https://doi.org/10.1161/HYPERTENSIONAHA.125.25305>
- 40 Mario Aguilar, “Cadence raises \$100 million to automate chronic disease care with regulated AI,” *STAT*, June 23, 2026.
<https://www.statnews.com/2026/06/23/cadence-100-million-raise-automate-chronic-care-ai/>
- 41 U.S. Food and Drug Administration (FDA), “Participants Selected for TEMPO for Digital Health Devices Pilot,” last modified August 21, 2026.
<https://www.fda.gov/medical-devices/digital-health-center-excellence/participants-selected-tempo-digital-health-devices-pilot>

References

- ⁴² Richard J. McManus, Jonathan Mant, Emma P Bray, et al., “Telemonitoring and Self-Management in the Control of Hypertension (TASMINH2): A Randomised Controlled Trial,” *The Lancet* 376, no. 9736 (2010): 163–172.
[https://doi.org/10.1016/S0140-6736\(10\)60964-6](https://doi.org/10.1016/S0140-6736(10)60964-6)
- ⁴³ Richard J. McManus, Jonathan Mant, Marloes Franssen, et al., “Efficacy of Self-Monitored Blood Pressure, with or without Telemonitoring, for Titration of Antihypertensive Medication (TASMINH4): An Unmasked Randomised Controlled Trial,” *The Lancet* 391, no. 10124 (2018): 949–959.
[https://doi.org/10.1016/S0140-6736\(18\)30309-X](https://doi.org/10.1016/S0140-6736(18)30309-X)
- ⁴⁴ Karen L. Margolis, Stephen E. Asche, Anna R. Bergdall, et al., “Effect of Home Blood Pressure Telemonitoring and Pharmacist Management on Blood Pressure Control: A Cluster Randomized Clinical Trial,” *JAMA* 310, no. 1 (2013): 46–56.
<https://doi.org/10.1001/jama.2013.6549>
- ⁴⁵ Biobeat, “First-Ever: FDA Clears Biobeat’s Wearable Watch and Patch for Non-Invasive Cuffless Monitoring of Blood Pressure,” press release, *PR Newswire*, August 26, 2019.
<https://www.prnewswire.com/ll/news-releases/first-ever-fda-clears-biobeats-wearable-watch-and-patch-for-non-invasive-cuffless-monitoring-of-blood-pressure-300906765.html>
- ⁴⁶ OMRON Healthcare, “OMRON Unveils a New Interactive Blood Pressure Monitoring Skill for Amazon Alexa,” January 30, 2019.
<https://www.omron-healthcare.com/health-and-lifestyle/omron-unveils-a-new-interactive-blood-pressure-monitoring-skill-for-amazon-alexa>
- ⁴⁷ Karen L. Margolis, Anna R. Bergdall, A. Lauren Crain, et al., “Comparing Pharmacist-Led Telehealth Care and Clinic-Based Care for Uncontrolled High Blood Pressure: The Hyperlink 3 Pragmatic Cluster-Randomized Trial,” *Hypertension* 79, no. 12 (2022): 2708–2720.
<https://doi.org/10.1161/HYPERTENSIONAHA.122.19816>
- ⁴⁸ Hunter Nichols, Christopher P. Cannon, Benjamin M. Scirica, and Naomi D. L. Fisher, “A Remote Hypertension Management Program Clinical Algorithm,” *Clinical Cardiology* 45, no. 12 (2022): 1147–1162.
<https://doi.org/10.1002/clc.23919>
- ⁴⁹ Anh T. L. Truong, Shi-Bei Tan, Golda Z. Wang, et al., “CURATE. AI-Assisted Dose Titration for Anti-Hypertensive Personalized Therapy: Study Protocol for a Multi-Arm, Randomized, Pilot Feasibility Trial Using CURATE.AI (CURATE.AI ADAPT Trial),” *European Heart Journal - Digital Health* 5, no. 1 (2024): 41–49.
<https://doi.org/10.1093/ehjdh/ztd063>
- ⁵⁰ Omada Health, “Omada Health Launches Physician-Guided Care Program, the First Virtual Cardiometabolic Clinic,” April 14, 2021.
<https://www.omadahealth.com/resource-center/omada-health-launches-physician-guided-care-program-the-first-virtual-cardiometabolic-clinic>
- ⁵¹ Caitlyn Durcan, “Introducing Cadence, a New Health Tech Company Redefining Remote Care for the Nation’s Largest Health Systems,” *Cadence*, August 5, 2021.
<https://www.cadence.care/post/introducing-cadence-new-health-tech-company>
- ⁵² Collier et al., “Personalized Antihypertensive Treatment,” e030749.
- ⁵³ Yi et al., “Development of a Machine Learning-Based Model,” 1660–1669.
- ⁵⁴ Wang et al., “Large Language Model Agent,” 212–224.
- ⁵⁵ FDA, “510(k) Premarket Notification Substantial Equivalence Determination Letter: GO Blood Pressure Monitoring System (GO),” July 2, 2025.
https://www.accessdata.fda.gov/cdrh_docs/pdf25/K250415.pdf
- ⁵⁶ Mario Aguilar, “Cadence raises \$100 million.”
- ⁵⁷ FDA, “Participants Selected for TEMPO.”
- ⁵⁸ National Institute of Mental Health, “Mental Illness,” last updated September 2024.
<https://www.nimh.nih.gov/health/statistics/mental-illness>
- ⁵⁹ National Center for Health Workforce Analysis, “State of the Behavioral Health Workforce, 2025,” Health Resources & Services Administration, Bureau of Health Workforce, December 2025.
<https://bhw.hrsa.gov/sites/default/files/bureau-health-workforce/data-research/Behavioral-Health-Workforce-Brief-2025.pdf>
- ⁶⁰ Max Rollwage, Jessica McFadyen, Keno Juchems, et al., “A Cognitive Layer Architecture to Support Large-Language Model Performance in Psychotherapy Interactions,” *Nature Medicine* 32, no. 5 (2026): 1717–1725.
<https://doi.org/10.1038/s41591-026-04278-w>
- ⁶¹ Myra Cheng, Cinoo Lee, Pranav Khadpe, et al., “Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence,” arXiv (2025): 2510.01395.
<https://doi.org/10.48550/arXiv.2510.01395>
- ⁶² Kathleen Kara Fitzpatrick, Alison Darcy, and Molly Vierhile, “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 Mental Health* 4, no. 2 (2017): e19.
<https://doi.org/10.2196/mental.7785>
- ⁶³ Becky Inkster, Shubhankar Sarda, and Vinod Subramanian, “An Empathy-Driven, Conversational Artificial Intelligence Agent (Wysa) for Digital Mental Well-Being: Real-World Data Evaluation Mixed-Methods Study,” *JMIR mHealth and uHealth* 6, no. 11 (2018): e12106.
<https://doi.org/10.2196/12106>
- ⁶⁴ Fitzpatrick et al., “Delivering Cognitive Behavior Therapy,” e19.
- ⁶⁵ Inkster et al., “An Empathy-Driven,” e12106.

References

- ⁶⁶ Han Li, Yi-Chieh Lee, Robert E. Kraut, and David C. Mohr, “Systematic Review and Meta-Analysis of AI-Based Conversational Agents for Promoting Mental Health and Well-Being,” *npj Digital Medicine* 6, no. 1 (2023): 236. <https://doi.org/10.1038/s41746-023-00979-5>
- ⁶⁷ Adriana Krasniansky, Bill Evans, and Megan Zweig, “2021 Year-End Digital Health Funding: Seismic Shifts beneath the Surface,” Rock Health, January 10, 2022. <https://rockhealth.com/insights/2021-year-end-digital-health-funding-seismic-shifts-beneath-the-surface/>
- ⁶⁸ Heinz et al., “Randomized Trial of a Generative AI Chatbot,” Aloa2400802.
- ⁶⁹ Rollwage et al., “A Cognitive Layer Architecture,” 1717–1725.
- ⁷⁰ STAT, “Why Woebot, a Pioneering Therapy Chatbot, Shut Down,” July 2, 2025. <https://www.statnews.com/2025/07/02/woebot-therapy-chatbot-shuts-down-founder-says-ai-moving-faster-than-regulators/>
- ⁷¹ Elizabeth C. Stade, Zoe M. Tait, Samuel T. Campione, et al., “Real-World Use of Large Language Models for Mental Health in 2024,” *npj Digital Medicine* 9, no. 1 (2026): 630. <https://doi.org/10.1038/s41746-026-02842-9>
- ⁷² Slingshot AI, “Slingshot Launches Ash, the First AI Designed for Therapy,” press release, *Business Wire*, July 22, 2025. <https://www.businesswire.com/news/home/20250722566346/en/Slingshot-Launches-Ash-the-First-AI-Designed-for-Therapy>
- ⁷³ Lyra Health, “Lyra Health Introduces First Clinical-Grade AI for Mental Health,” press release, October 14, 2025. <https://www.lyrahealth.com/announcement/lyra-health-introduces-first-clinical-grade-ai-for-mental-health/>
- ⁷⁴ Ron C. Li, Leah Rosengaus, Lisa Gohil, et al., “A Framework for Using AI to Drive Care Model Transformation: Building Cars Rather Than Faster Horses,” *npj Health Systems* 3, no. 1 (2026): 7. <https://doi.org/10.1038/s44401-026-00071-6>
- ⁷⁵ Chris Miller-Rosales, Nancy E. Morden, Mary F. Brunette, et al., “Provision of Digital Health Technologies for Opioid Use Disorder Treatment by US Health Care Organizations,” *JAMA Network Open* 6, no. 7 (2023): e2323741. <https://doi.org/10.1001/jamanetworkopen.2023.23741>
- ⁷⁶ PHTI, “Virtual Opioid Use Disorder Solutions,” September 16, 2025. <https://phti.org/assessment/virtual-opioid-use-disorder-solutions/>
- ⁷⁷ Emanuel et al., “Will Autonomous AI Exceed AI-Aided Physicians,” 915–918.
- ⁷⁸ Dan Zeltzer, Zehavi Kugler, Lior Hayat, et al., “Comparison of Initial Artificial Intelligence (AI) and Final Physician Recommendations in AI-Assisted Virtual Urgent Care Visits,” *Annals of Internal Medicine* 178 (2025): 498–506. <https://doi.org/10.7326/ANNALS-24-03283>
- ⁷⁹ Ashwin Nayak, Sharif Vakili, Kristen Nayak, et al., “Use of Voice-Based Conversational Artificial Intelligence for Basal Insulin Prescription Management Among Patients With Type 2 Diabetes: A Randomized Clinical Trial,” *JAMA Network Open* 6, no. 12 (2023): e2340232. <https://doi.org/10.1001/jamanetworkopen.2023.40232>
- ⁸⁰ Utah Department of Commerce, “Utah and Doctronic Announce Groundbreaking Partnership for AI Prescription Medication Renewals,” news release, January 6, 2026. <https://commerce.utah.gov/2026/01/06/news-release-utah-and-doctronic-announce-groundbreaking-partnership-for-ai-prescription-medication-renewals/>
- ⁸¹ Alice Schwarze and Zach Boyd, “Evidentiary Expectations for Healthcare AI Sandbox: Demonstrating Patient-Focused, Safety-Driven AI Applications,” Utah Department of Commerce, Office of Artificial Intelligence Policy, August 2026. <https://commerce.utah.gov/wp-content/uploads/2026/08/Evidentiary-Expectations-for-Healthcare-AI-Sandbox.pdf>
- ⁸² State of Delaware, “Delaware Launches Bold AI Sandbox Initiative, Cementing Its Role as a National Leader in Responsible Tech Innovation,” news release, July 23, 2025. <https://news.delaware.gov/2025/07/23/delaware-launches-bold-ai-sandbox-initiative-cementing-its-role-as-a-national-leader-in-responsible-tech-innovation/>
- ⁸³ State of Connecticut, “Governor Lamont Signs Legislation Establishing Youth Online Safety Protections, Regulations Over Artificial Intelligence, and Initiatives to Upskill Connecticut’s Workforce,” news release, June 2, 2026. <https://portal.ct.gov/governor/news/press-releases/2026/06-2026/governor-lamont-signs-legislation-establishing-youth-online-safety-protections>
- ⁸⁴ FDA, “TEMPO for Digital Health Devices Pilot,” last modified July 22, 2026. <https://www.fda.gov/medical-devices/digital-health-center-excellence/tempo-digital-health-devices-pilot>
- ⁸⁵ ARPA-H, “ARPA-H launches the world’s first bid to build an FDA-authorized clinical AI for cardiovascular care,” news release, September 9, 2026. <https://arpa-h.gov/news-and-events/arpa-h-launches-worlds-first-bid-build-fda-authorized-clinical-ai-cardiovascular>
- ⁸⁶ FDA, “Considerations for the Regulation of Generative AI-Enabled Medical Devices: Discussion Paper and Request for Feedback,” last modified August 18, 2026. <https://www.fda.gov/medical-devices/digital-health-center-excellence/considerations-regulation-generative-ai-enabled-medical-devices-discussion-paper-and-request>
- ⁸⁷ Federation of State Medical Boards, “Federation of State Medical Boards Announces New Workgroup on Regulation of Artificial Intelligence in Medical Practice,” news release, May 19, 2026. <https://www.fsmb.org/advocacy/news-releases/fsmb-announces-new-workgroup-on-regulation-of-artificial-intelligence-in-medical-practice/>

References

- ⁸⁸ Michelle M. Mello and Neel Guha, “Understanding Liability Risk from Using Health Care Artificial Intelligence Tools,” *New England Journal of Medicine* 390, no. 3 (2024): 271–278.
<https://doi.org/10.1056/NEJMhle2308901>
- ⁸⁹ CMS, “ACCESS (Advancing Chronic Care with Effective, Scalable Solutions) Model,” last modified September 18, 2026.
<https://www.cms.gov/priorities/innovation/innovation-models/access>
- ⁹⁰ CMS, “CMS Adds Medicare Options to Help People Manage Substance Use Disorder, Heart Failure, and Other Common Chronic Conditions,” press release, September 15, 2026.
<https://www.cms.gov/newsroom/press-releases/cms-adds-medicare-options-help-people-manage-substance-use-disorder-heart-failure-other-common>
- ⁹¹ Office of the National Coordinator for Health Information Technology (ONC), “ONC Finalizes the Adoption of Certain Health IT Standards in the FY2027 CMS IPPS Final Rule,” fact sheet, July 2026.
<https://healthit.gov/resources/onc-finalizes-the-adoption-of-certain-health-it-standards-in-the-fy2027-cms-ippss-final-rule/>
- ⁹² CMS, “Innovation Insight: Major Health Plans Join ACCESS Payer Pledge,” February 12, 2026.
<https://www.cms.gov/priorities/innovation/innovation-insight-major-health-plans-join-access-payer-pledge>

