Healthcasts

The Human Algorithm: Confronting the Role of AI in Healthcare

Why the future of healthcare is human insight + AI, and how we should be amplifying humanity in healthcare through technology, not replacing it.

Table of Contents

03 Executive Summary

04 Challenges in AI Adoption

06 Where AI in Healthcare Stands Today

09 Why Human Judgment Still Matters

10 Navigating the Gray Areas of Care

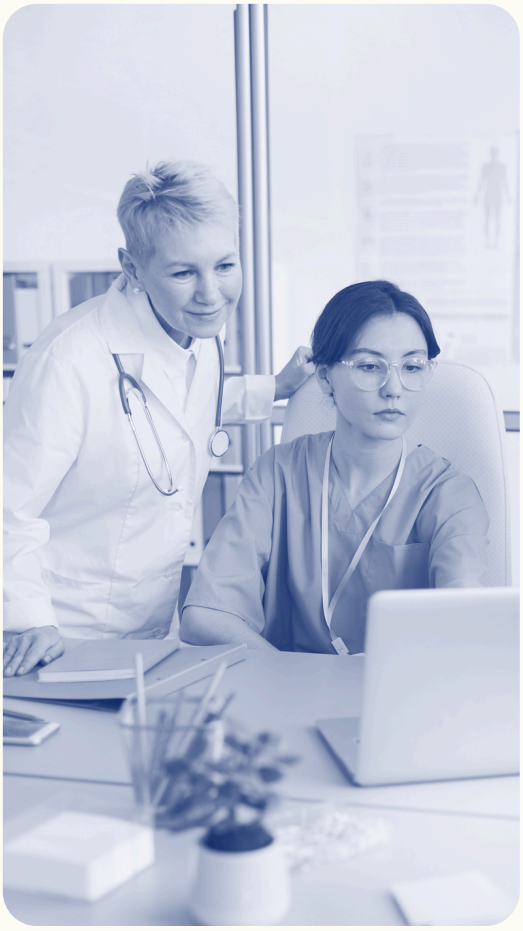
14 Healthcasts' Vision for the Next Decade of Healthcare

15 Redefining Intelligence in Medicine

Executive Summary



2 in 3 physicians are now utilizing AI



At Healthcasts, we believe that the future of healthcare won't be determined by the capabilities of AI, but by the human layer wrapped around it — empathy, judgment, and nuance.

According to an AMA survey, 2 in 3 physicians are now utilizing AI, yet 47% cited increased oversight as the number one regulatory action needed to increase their trust in adoption.[1]

While the utilization of AI is clearly accelerating, the data shows trust depends on more than technology —it also requires credibility that only humans can provide.

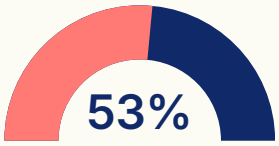
This white paper explores how we can build that future: the key barriers, the enablers, what is shifting today, and how physicians can act now.

Our core belief:

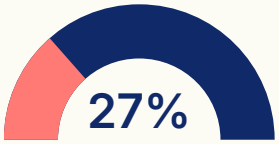
The future of healthcare is human insight + AI.

Challenges in AI Adoption

THE TRUST GAP



of physicians on Healthcasts have previously chosen not to follow an AI recommendation.



agreed that overconfidence and hallucinations are the greatest limitations of current AI tools in clinical care, according to physicians on Healthcasts.

We're living in a moment where the volume of available information has never been higher, and as a result, valuable insights are buried under more noise than ever, leaving clinicians to separate what's credible from what's simply loud.

Traditionally reliable tools like Google, which once prioritized research papers, authoritative clinical sites, and third-party surveys, increasingly surface opinions from unverified sources like Reddit threads or Quora responses. These platforms can be valuable for many types of questions where subjective perspectives pose little risk. *Which running shoes have the best cushioning? What's a great recipe for a weeknight dinner?*

But in healthcare? Accuracy isn't just helpful—it's foundational. For clinical decision-making, the provenance of information matters as much as the information itself. And this is precisely where concern around AI's role in care is growing.

Despite enormous progress in recent years, AI-generated answers still struggle with reliability. *The New York Times* recently reported that AI systems "do not—and cannot—decide what is true and what is false." They generate plausible text, not verified facts. In some tests, hallucination rates of newer systems reached as high as 79%—meaning the majority of answers contained fabricated or distorted information.[2]

This creates a real-world consequence: even when clinicians find AI helpful, they do not always trust it. Does this mean AI should be removed from clinical decision-making? Absolutely not. But it does mean AI cannot operate alone. Clinicians need to provide the context and nuance necessary to vet AI-generated insights appropriately.

THE BLACK BOX PROBLEM

You may remember the phrase "show your work" from elementary school math lessons, and in medicine, it's not just a suggestion—it's a mandate.

When someone's health is on the line, clinicians must go beyond a simple recommendation to provide the real "**why**" behind their decision. Patients (rightfully) want to know the reasoning for a new treatment or diagnostic test, and how it fits their specific needs. This is where the use of AI for clinical purposes gets tricky. It is incredibly helpful in selecting and summarizing content across the internet, but its ability to show how it came to its conclusion is limited to simply linking to sources.



Imagine this:

A clinician is treating a patient newly diagnosed with type 2 diabetes. They turn to an AI tool for guidance on first-line treatment options. In seconds, the AI aggregates evidence-based recommendations from trusted sources and offers a concise summary:

"Metformin is generally considered the first-line therapy for type 2 diabetes. Here are links to ADA guidelines and recent clinical studies supporting this approach."

This is undeniably useful. It saves time, pulls relevant evidence, and offers a solid starting point. But it stops short of what real clinical care requires.

AI can't tell you:

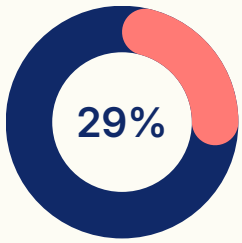
- Whether this patient's kidney function makes metformin appropriate
- How their past medication experiences influence adherence
- Whether lifestyle factors or comorbidities warrant an alternative
- Or how shared decision-making might shape the choice

AI offers information. Only clinicians can offer interpretation.

AI may recommend an option based on patterns in the data, but it cannot meaningfully articulate the reasoning or risk-benefit analysis behind it. Clinicians must still evaluate the proposal, weigh its appropriateness, and justify the path forward.

In the end, AI can surface recommendations. It can accelerate the research. It can streamline the workload. However, it's the human layer—critical thinking, professional judgment, and clinical context—that transforms raw data into trustworthy care.

Where AI in Healthcare Stands Today



of physicians on Healthcasts said AI informs at least 10% or more of their decisions.



If we want to understand how AI and humans can create better outcomes together moving forward, we must first understand the current state of this relationship.

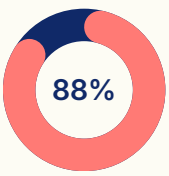
The future, and most importantly, the success of healthcare, cannot simply hinge on the race to efficiency.

If AI becomes just a tool for speed, we've misunderstood its true potential. AI has shown enormous promise, but reports reveal there is still hesitation and a strong desire for human "guardrails."

- A 2023 nationally representative survey (2,039 US adults) found that 65.8% of respondents expressed low trust in their healthcare system's ability to use AI responsibly, and 57.7% lacked confidence that their healthcare systems would prevent AI tools from causing harm.[3]

University of Minnesota School of Public Health

- Another study showed that nearly 9 in 10 American adults (88%) express concerns about the increasing use of AI in healthcare settings, with 58% worried about the lack of human oversight.[4]





express concerns about the increasing use of AI in healthcare settings

Gwynedd Mercy University

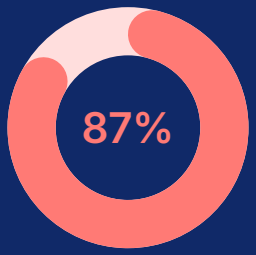
Additionally, the Medscape & HIMSS AI Adoption in Healthcare Report 2024 revealed that 86% of physicians already leverage AI in their medical organizations, with 60% recognizing its ability to uncover health patterns and diagnoses beyond human detection.[5] However, concerns persist, with 72% citing data privacy as a significant risk.[5]

But there are bright spots for the success of AI's use in medicine, especially when working in tandem with humans.

- According to a study published in the *European Journal of Cardiovascular Medicine*, integrating AI tools into diagnostic processes dropped the rate of diagnostic errors from 22% to 12%. Predictive algorithms are identifying patients at risk for sepsis hours before clinical symptoms appear.[6] Natural language models are helping doctors complete documentation faster, freeing time for patient care. The challenge isn't whether AI can deliver results — it's ensuring it does so with human oversight and accountability at its core.

- In 2023, a study in *The Lancet* found that AI-assisted mammography in Sweden reduced radiologist workload by at least 44% while catching as many or more cancers than standard reads. The key was collaboration: AI flagged suspicious scans, but radiologists made the final call.[7]

HEALTHCASTS' SURVEY AND RESULTS



of physicians on Healthcasts agree or strongly agree that AI can provide information, but it cannot replace the judgment that comes from human experience.

Our core belief is simple: to successfully utilize AI in healthcare, we must ensure it amplifies clinicians' expertise, not replace it. To explore how this plays out in practice, we surveyed our community members on their sentiments toward AI, including how they're using these tools and where human judgment is being applied alongside them.

What we found wasn't surprising: physicians showed clear hesitancy about integrating AI deeper into their clinical workflow. However, many also expressed optimism about its potential to enhance patient care.

Much of the debate around AI in clinical decisions centers on which tasks it's most helpful for. Within our community, its top three uses reported by physicians are developing treatment plans, followed closely by administrative tasks or other tasks, then interpreting lab or imaging results.

As physicians incorporate AI into their workflows, one point remains clear: human judgment is indispensable. That belief came through strongly in our survey results, with 87% of respondents saying they agree or strongly agree that AI can provide information, but it cannot replace the judgment that comes from human experience.

When asked how often they felt that reliance on AI could lead to mistakes in clinical settings if not carefully supervised, over three-quarters (77%) selected "sometimes" or "often", with 6% saying "always", signaling a clear preference for human judgment to make the final call, and reinforcing that the best results happen when clinical expertise and AI work together.

Yet, their confidence in using AI-generated results to make clinical decisions was promising:



20% expressed they were not confident in using AI-generated recommendations to make clinical decisions

37% showed slight confidence

31% reported moderate confidence

12% had high confidence (very or extremely confident)



Within the Healthcasts community, when asked what percentage of their clinical decisions are currently supported by AI, **32%** of physicians reported "none", while **40%** indicated that AI informs fewer than **10%** of their decisions. However, **29%** said AI informs at least **10%** of their decisions—falling within ranges of **10–25%, 26–50%**, or even more than **50%**. The road to mass adoption is slow, but it shows plenty of hope for AI playing a supporting role for physicians.

All of the findings above, while demonstrating a clear divide between skepticism and the undeniable signs of progress, tell a consistent story: the effectiveness of AI ultimately depends on the human judgment guiding it.

Why Human Judgment Still Matters

Evidence continues to show that when machines and humans work together, outcomes improve—not because technology replaces judgment, but because it enhances it.

The most current evidence shows that hybrid human-AI collectives (teams composed of both human experts and machine learning systems) achieve superior outcomes by leveraging the strengths of each. A recent medical study demonstrated this concept, finding that diagnostic collectives of physicians and large language models (LLMs) were significantly more accurate than either human-only or AI-only teams—specifically, where LLMs missed the correct diagnosis, individual physicians often contributed the correct diagnosis. While data can accelerate diagnostics, human discernment is key to accuracy.[8]

That's the human layer in action. Even when an AI tool provides citations, those references can still be fabricated, outdated, or misinterpreted. No model, not even one trained specifically for healthcare, is immune to these pitfalls, which means the **clinician's role shifts from simply retrieving information to interrogating it.**



As one *Family Medicine* article notes, even tools like Open Evidence, while powerful at the point of care, can create a risk of "early closure" for novice or fatigued clinicians because of their brevity and narrow focus.[9] Without deeper exploration, it's easy to accept concise answers without questioning what might be missing.

The lesson? **Machines can accelerate discovery, but only people can ensure reliability.** In an era defined by abundant data and instant answers, critical thinking becomes the clinician's most valuable competency—the ability to evaluate, interpret, and contextualize information before acting on it. It's not enough to ask what the data says; we must ask whether it makes sense, whether it fits the patient, and whether it holds up under scrutiny.

Navigating the Gray Areas of Care

Few areas reveal the limits of AI more clearly than off-label prescribing—the practice of using an FDA-approved medication for a purpose other than what it was originally approved for. It's far more common than many realize. According to a study published in 2025 by the *Journal of Pharmaceutical Policy and Practice*, about **25% of prescription uses are off-label.**[10]

Think:

- Using Ozempic for weight loss instead of exclusively for type 2 diabetes.
- Prescribing antidepressants for chronic pain.
- Using anti-seizure medications to prevent migraines.

These decisions aren't reckless—they reflect the reality that medicine evolves faster than formal guidelines. Clinicians regularly adapt treatments to the patient in front of them, drawing on experience, emerging evidence, and conversations with trusted peers. Within the Healthcasts community itself, this nuance is clear: **46%** of practitioners prescribe off-label medication "sometimes", and **24%** do so "often".

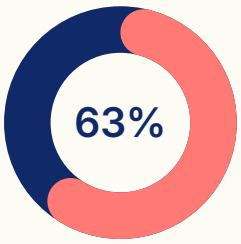
But this is precisely where AI struggles.

AI tools are built to excel in structured, codified knowledge such as guidelines, published studies, clinical summaries, and known pathways. Off-label care, by contrast, often lives in the gray space between the evidence and the exam room.

This is why peer conversations become indispensable. When guidelines don't exist, the only real gut check is another clinician who has seen similar patients and tried similar approaches. Off-label prescribing is often based on accumulated wisdom—the kind of tacit knowledge AI simply cannot replicate.

Off-label care is where medicine becomes art as much as science, and that is something only humans can do.

FINDING THE "AND" WITH AI



of practitioners on Healthcasts said they trust peer-to-peer discussion the most.



Not every case reflects a typical textbook scenario that standard guidelines can address. In fact, within the Healthcasts community, the most common response from 38% of practitioners was that 11 to 30% of their patients fall into "non-textbook" cases—potentially up to a third of their entire caseload.

When asked whether they trust AI insights, peer-to-peer discussion, or both equally for solving complex, non-textbook cases, 63% of practitioners said they trust peer-to-peer discussion the most, while 31% trust both equally.

The examples below demonstrate how combining AI's efficiency with human expertise leads to more informed decisions and better patient outcomes, particularly for unique cases.

Case Study A: Fever of Unknown Origin in a Returning Traveler

A 40-year-old man presents with prolonged fever, night sweats, and weight loss after travel to rural sub-Saharan Africa. Initial workup—including malaria smears, blood cultures, and imaging—is inconclusive. A practitioner on Healthcasts wanted to understand what differential diagnosis they should prioritize and what further diagnostics to run.

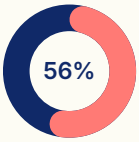
AI-generated differential and diagnostic:

"Priority differentials: Visceral leishmaniasis (kala-azar), acute HIV infection, brucellosis, typhoid fever (culture-negative), tuberculosis, and hematologic malignancy (eg, lymphoma).

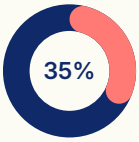
Advanced diagnostics: rK39 antigen or PCR for *Leishmania*, HIV antigen/antibody testing, *Brucella* serology, bone marrow biopsy (for *Leishmania*, TB, or malignancy), and repeat thick/thin malaria smears or PCR to rule out low-level parasitemia."

The Healthcasts Consensus:

*The Healthcasts Consensus utilizes AI to pull together all comments on a post and breaks down recommendations to show what physicians agree on for next steps.



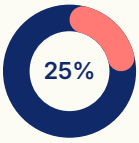
"Malaria remains a top priority diagnosis."



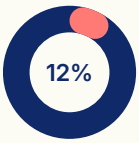
"Bacterial and viral etiologies (HIV, TB, typhoid, brucellosis)."



"Visceral leishmaniasis and other parasitic infections."



"Early and urgent infectious disease (ID) consultation advised."



"Non-infectious causes like malignancy/HLH considered last."

Clinicians also shared nuanced observations in their comments, including:

- Repeat malaria testing and broaden the infectious workup to include region-specific diseases such as visceral leishmaniasis and dengue.
- Gather detailed travel and exposure history, including insect bites, geographic regions visited, and potential sexual or GI exposures.
- Consult infectious disease specialists and reference CDC travel resources to guide advanced diagnostics and determine isolation precautions.



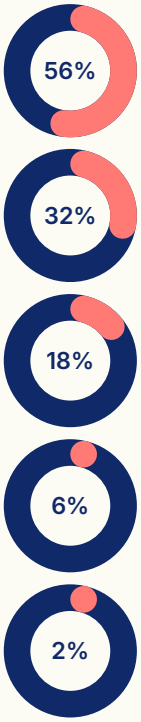
Case Study B: Adverse Side Effects in Patients Taking GLP-1s

GLP-1 receptor agonists are increasingly used for weight loss and diabetes management, but side effects vary widely. A practitioner on Healthcasts sought to understand the adverse effects practitioners were seeing in patients taking GLP-1 RAs, whether for diabetes or weight loss.

AI-generated answer:

Commonly reported adverse effects include nausea, vomiting, diarrhea, constipation, and bloating—typically dose-dependent and transient. Less frequent issues include delayed gastric emptying, gallbladder events, and fatigue; rare but serious effects include pancreatitis and bowel obstruction.

The Healthcasts Consensus:



"Gastrointestinal side effects are most common."

"Rare but severe risks: pancreatitis and gallbladder issues."

"Muscle loss and nutritional concerns observed."

"Mood changes and mental health effects noted."

"Medication cost and access affect treatment adherence."

Additionally, practitioners added clinical depth:

- Cyclic vomiting and persistent constipation were reported by their patients even months after discontinuation, suggesting some side effects may have longer-term impacts than commonly recognized.
- Hydration strategies and the use of antiemetics like ondansetron were cited as effective for managing nausea, especially during dose escalation.

- Higher incidence of GI side effects, particularly constipation and gastroparesis, was noted in non-diabetic patients, influencing decisions around patient selection and dosing.

- Adjusting injection intervals (eg, every 10 days instead of 7) appeared to reduce side effect severity in some cases, offering a potential strategy for improving tolerability.

These examples of AI alongside the responses found on Healthcasts reinforce our point: AI doesn't replace human expertise—it amplifies it to enhance treatment plans with the human knowledge of what's worked before. The key is collaboration, not competition.

Healthcasts' Vision for the Next Decade of Healthcare

For over 20 years, Healthcasts has connected practitioners with real-world clinical insights, and we continue to advance this mission in step with technology—not against it. We serve as a bridge between AI's data-driven promise and the human judgment, empathy, and nuance that make care meaningful. With our platform, practitioners are empowered to:

- Connect with a community of NPI-verified peers to share real-world clinical insights that go beyond guidelines.
- Post cases and questions to receive guidance from physicians who have navigated similar challenges.
- Access AI-powered Consensus summaries that pull peer recommendations together, showing what physicians agree on for next steps.
- Search across our library of peer insights and cases to find nuanced answers, fast.

In doing so, Healthcasts contributes directly to the realization of a healthcare ecosystem in which AI is powerful because it is grounded in human experience, trust, and relevance.

Redefining Intelligence in Medicine

Healthcare won't be defined by the smartest algorithms, the fastest systems, or the most data—it will be defined by how those algorithms, systems, and data are joined to humanity. If the next decade is built on trust, empathy, and human insight, then technology becomes not the story but the powerful support tool that uplifts patient care.

Healthcasts believes the future is not AI versus humans—it's AI with humans.
And that is where real progress lies.



Sources

1. Henry TA. 2 in 3 physicians are using health AI—up 78% from 2023. American Medical Association. February 26, 2025. Accessed December 9, 2025. https://www.ama-assn.org/ practice-management/digital-health/2-3-physicians-are-using-health-ai-78-2023

2. Metz C, Weise K. A.I. Is Getting More Powerful, but Its Hallucinations Are Getting Worse. New York Times. May 5, 2025. Accessed December 9, 2025. https://www.nytimes.com/2025/05/05 /technology/ai-hallucinations-chatgpt-google.html#.

3. McDill V. Most people do not trust healthcare systems to use artificial intelligence (AI) responsibly. University of Minnesota School of Public Health. March 6, 2025. Accessed December 9, 2025. https://www.sph.umn.edu/news/most-people-do-not-trust-healthcare-systems-to-use-artificial-intelligence-ai-responsibly/

4. Business Wire. More than Half of Americans Believe AI will Improve Healthcare Outcomes & Lower Costs, New Gwynedd Mercy University Survey Finds. January 15, 2025. Accessed December 9, 2025. https://www.businesswire.com/news/home/2025011502 9465/en/More-than-Half-of-Americans-Believe-AI-will-Improve-Healthcare-Outcomes-Lower-Costs-New-Gwynedd-Mercy-University-Survey-Finds

5. Medscape, HIMSS. Early Successes, Untapped Potential, Lingering Questions: AI ADOPTION IN HEALTHCARE REPORT 2024. https://cdn.sanity.io/files/sqo8bpt9/production/68216fa5d161adebceb50b7add5b496138a78cdb.pdf. Published December 6, 2024. Accessed December 9, 2025.

6. Borkar SR, De A. Artificial Intelligence in Internal Medicine: A Study on Reducing Diagnostic Errors and Enhancing Efficiency. Eur J Cardiovasc Med. 2025;15(9):105-111. doi:10.61336/ejcm/25-09-19.

7. Dembrower K, Crippa A, Colón E, Eklund M, Strand F, et al. Artificial intelligence for breast cancer detection in screening mammography in Sweden: a prospective, population-based, paired-reader, non-inferiority study. Lancet. 2023;5(10):e703-e711.

8. Zöller N, Berger J, Lin I, Fu N, Komarneni J, et al. Human–AI collectives most accurately diagnose clinical vignettes, Proc Natl Acad Sci. 2025;122(24):e2426153122. https://doi.org/10.1073/pnas.2426153122

9. Wu V, Casauay J. OpenEvidence. Fam Med. 2025;57(3):232-233. DOI: 10.22454/FamMed.20 24.587513

10. Blankart KE, Lichtenberg FR. Prevalence and relationship with health of off-label and contraindicated drug use in the United States: a cross-sectional study. J Pharm Policy Pract. 2025;18(1):2472221. Published March 6, 2025. doi:10.1080/20523211.2025.2472221.

Healthcasts

healthcasts.com hello@healthcasts.com 775-278-6325