

Agentic Medical Information and the New Operating Standard



AI cannot be trusted as the ultimate authority on clinical matters. Yet it is quickly becoming the first place healthcare professionals (HCPs) and patients turn to search for and access information. This creates an important inflection point for medical information. How should the function redefine its role within the industry, and what must change about the way it operates to succeed in this new environment?

In this paper, we explore this new information era and why now is the time for medical information to double down on its role as custodian of scientific truth and orchestrate the seamless experiences customers now expect.



HCPs and patients shift to AI-first information seeking

In the wild west of information, where limitless content competes for mindshare on any given topic, it's difficult to cut through the noise. AI is transformative for how it solves exactly this problem, surfacing the information users need in seconds.

How AI elevates information seeking

- | | |
|---------------------------------------|--|
| ✓ On-demand immediacy | Delivers instantaneous responses when prompted |
| ✓ Frictionless interactions | Replaces scattered, multi-site navigation with a single conversational search interface |
| ✓ Seamless ecosystems | Maintains continuity across conversations, devices, and applications |
| ✓ Hyperpersonalization | Adapts responses based on user history, preferences, and custom instructions |
| ✓ Anticipation and proactivity | Suggests follow-up questions, related actions, and recommendations |
| ✓ Conversational fluency | Enables human-like dialogue that feels collaborative, with users able to define tone, expertise, or role |

By removing the friction that has traditionally defined research, AI has fundamentally changed how people seek, consume, and act on information. Google AI Overviews have accelerated the rise of zero-click search, where users receive a summarized answer without ever visiting the original source. Today, 68% of Google searches end without a click to an external website.¹ AI Overviews appear on 88% of healthcare searches and about 75% of biopharma-related keywords now trigger an AI-generated summary.²

The zero-click search phenomenon

68%

of Google searches now end without a single click to an external website

75%

of biopharma keywords trigger an AI Overview

88%

AI Overview trigger rate for healthcare related queries leads all industries

Among HCPs, the clinical use of AI continues to grow. Over 80% of physicians use AI professionally, doubling since 2023.³ As adoption increases, so does confidence, with 75% of physicians believing AI improves their ability to care for patients. Purpose-built clinical AI platforms such as OpenEvidence have seen unprecedented adoption, hitting over 1 million consultations in a single day at the point of care.⁴

HCP AI adoption

+80%

of physicians report using AI in a professional context

75%

of physicians believe AI improves their ability to care for patients

1M

consultations in a single day on OpenEvidence

¹ "In 2026, Less than One Third of Google Searches Still Send a Click," SparkToro, June 2026

² "18 Months of AI Overviews," BrightEdge, January 2026; "2026 Pharma AI Overviews Report," EVERIANA, May 2026

³ "2026 Physician Survey on Augmented Intelligence," American Medical Association, Center for Digital Health and AI, March 2026

⁴ "OpenEvidence Achieves Historic Milestone," PR Newswire, March 2026



AI model audit exposes critical information gap

Despite mainstream use among HCPs and patients, public large language models (LLMs) rely on a fundamentally flawed information environment. They are trained on accessible public data, lacking access to proprietary, evidence-based research stewarded by biopharmas. LLMs often stitch references together from secondary texts, such as Wikipedia footnotes rather than primary medical texts, often because a significant portion of medical literature is locked behind academic publishing paywalls.


Research into public AI tools has exposed a critical information gap. These tools are becoming the unvalidated front door to medical information, not only for HCPs, but for patients and caregivers seeking answers before or after clinical visits. The problem is that nearly 50% of medical responses are problematic or inaccurate, delivered with false confidence and hallucinated references.⁵ This leads to real world consequences at the point of care: erosion of trust, non-evidence-based treatment, and strains on healthcare systems and providers.

The clinical coin flip: key findings

[A recent study](#) evaluated five leading public AI models — Gemini, ChatGPT, DeepSeek, Meta AI, and Grok — using a series of closed- and open-ended prompts across health topics known to be vulnerable to misinformation. The prompts reflected the types of questions consumers and HCPs commonly ask AI, while responses were evaluated for clinical accuracy, readability, and the quality of supporting references.

Problematic and inaccurate AI medical responses

Nearly
50%
of AI medical responses
are problematic or
inaccurate



AI tools answered medical questions
with high confidence and certainty
99% of the time, even when
hallucinating or fabricating data

⁵ "Generative artificial intelligence-driven chatbots and medical misinformation: an accuracy, referencing and readability audit," BMJ Open, April 2026

Accuracy varies across AI models

Model	Rate of problematic response	Key insight
Gemini	40%	Produced fewest “highly problematic” errors
DeepSeek	48%	Performed well on formatting details
MetaAI	50%	Only model to trigger formal medical “refusal to answer”
ChatGPT	52%	Consistently used highest word-per-sentence complexity
Grok	58%	Produced an excess of “highly problematic” content (30%)

Performance depends on the topic



Stronger performing categories included fields characterized by high-quality, heavily structured literature and universally reinforced foundational concepts (i.e., vaccines and cancer)



Weaker performing categories included fields more subject to online misinformation, conflicting commercial interest, and unregulated clinical marketing (i.e., stem cells, athletic performance, and nutrition)

Referencing remains critical weakness

Not a single AI model was capable of producing a fully accurate, complete, and authentic reference list. When tasked with providing 10 scientific peer-reviewed citations to substantiate their closed-ended answers, the models faced serious limitations:

- AI tools frequently suffered from hallucinations such as inventing authors, mismatching publication years, and fabricating DOIs.
- ChatGPT explicitly acknowledged fabrication, stating that it fabricates information to maintain the appearance of completeness, even if that means sacrificing accuracy.

The implications of weak referencing extend beyond inaccurate citations. Medical information teams operate under rigorous standards of evidence validation and expert review when synthesizing literature for responses. Public AI tools, by contrast, generate responses based on patterns in their training data rather than a built-in understanding of clinical truth. Because the majority of peer-reviewed medical literature is inaccessible behind publishing paywalls, AI-generated responses are often based on secondary interpretations of data rather than primary evidence. When those responses are delivered in a human-like, authoritative tone, few think to question the validity.

False credibility poses high risk

LLMs are designed to communicate like a peer, prioritizing conversational fluency over scientific validation. For users asking health-related questions, human-like answers can create a false sense of authority. AI responses that sound technically detailed and confident are easy to trust at face value. As a result, the spread of medical misinformation and disinformation amplifies, increasing the risk of harmful health decisions and adverse public health outcomes.

Leading global organizations have sounded the alarm on the real-world consequences of medical misinformation and disinformation.



The [World Health Organization](#) has warned that health misinformation and disinformation is a major threat to global health.



The [World Economic Forum](#) cites misinformation and disinformation as the top short-to-medium term global risk.



The [European Parliament](#) identifies health disinformation as a systemic, ongoing threat to public health and democracy.



[ECRI](#) cites wide availability and viral spread of medical misinformation as the top 3 patient safety concern.

The trust gap between patients and providers is widening as a result. Industry surveys report that 67% of patients believe they have been exposed to misinformation about their condition in the past year, and 61% of HCPs explicitly state that misinformation has caused a loss of trust between them and their patients.⁶

⁶ "Doctored Truths urges pharma to lead the fight against health misinformation," Havas Lynx, October 2025




Medical information's critical inflection point

The reality of AI necessitates a shift: medical information must evolve from being a reactive gatekeeper to an active orchestrator of scientific knowledge. It's critical for medical information to meet HCPs where they are, embedding trusted scientific content into the digital channels they already use rather than expecting them to navigate traditional biopharma channels.

There is a disconnect between how the industry invests in engagement and how HCPs prefer to consume information. Biopharmas overinvest in traditional channels such as face-to-face engagement while digitally-savvy HCPs value searchable information that's accessible 24/7. In fact, **73% of HCPs** prefer virtual details about product topics, such as efficacy and safety and **75% of HCPs** would like to keep or increase their digital interactions with biopharmas. In an AI-first world, accessibility isn't just a customer experience imperative, it's essential to ensuring accurate medical information reaches clinicians before misinformation does.

The Evolution: Medical Information in 2030

The journey toward 2030 requires medical information leaders to assess where they have been, where they are, and where they must go.





The new agentic operating standard

Becoming the orchestrator of scientific knowledge requires rethinking the medical information operating model. Human expertise remains essential but the nature of work fundamentally changes. The future operating standard is deeply connected, human-centric, and agentic, empowering experts with intelligent technology that works alongside them and modern strategies that bridge internal capabilities with external value.

Embrace modern strategies to improve external accessibility

As AI reshapes how scientific information is discovered and consumed, medical information teams will need to evolve how they author and publish content. Success requires embracing modern strategies like structured content and Generative Engine Optimization (GEO) to ensure trusted scientific information is discoverable across digital channels and within AI-powered experiences.

This shift means moving beyond publishing static documents to managing machine-readable science, where clinically-validated knowledge can be dynamically assembled, surfaced, and accurately cited by LLMs. FAQs and Standard Responses are a high-value starting point for medical information teams. Authoring them as structured reusable components that can be managed and tagged like data enables conversational AI to surface them instantly at the point of need.

Elevate human expertise with agentic AI

Medical information teams are composed of highly trained experts with deep scientific knowledge. Yet many teams still spend valuable time navigating disconnected systems, manually entering data, and doing administrative work that adds little scientific value. Operational transformation is not for efficiency's sake alone — it is key to delivering a superior customer experience. Technology's role in this environment is to remove friction so experts can prioritize engaging healthcare professionals and advancing scientific understanding.

In an agentic medical information ecosystem, AI automates manual, repetitive work while humans focus on high-value clinical execution. As keepers of rich, proprietary data that HCPs cannot access using general LLMs, human specialists can spend their time conducting the deep research needed to answer more complex, nuanced questions. Scientific judgment, clinical reasoning, empathy, and relationship building remain uniquely human capabilities. Technology simply makes those capabilities easier to apply at scale.

The agentic medical information workflow in practice

The following four use cases illustrate how AI and human expertise work together across the agentic medical information workflow.



Case processing

Intelligent agents sit at the front door, managing intake, triage, routing, and documentation, automatically capturing required information and completing compliance-driven administrative tasks within the workflow before a specialist engages the case.



Response enablement

Embedded AI drafts responses using approved structured content, standard responses, and scientific references. Rather than creating new content, agents rapidly compile trusted information for expert review, accelerating response times while maintaining scientific integrity with human validation.



Research assistance

AI serves as an always-available scientific assistant, surfacing relevant literature, prior interactions, internal guidance, and supporting evidence within the user's existing workspace. Experts receive the information they need without leaving the platform, allowing them to focus entirely on clinical interpretation and meaningful scientific dialogue.



Insight generation

Connected agents continuously analyze inquiries, interactions, and enterprise data to identify emerging trends, educational gaps, and evolving information needs. Instead of simply reporting historical activity, medical information becomes a proactive intelligence function that helps shape medical strategy across the organization.



Leading the next era of medical information

As AI-first information seeking becomes the default for healthcare-related matters, medical information must work to bridge the critical information gap between scientific rigor and external accessibility. Doing so requires examining internal capabilities and how they can evolve to support new realities and meet the liquid expectations of HCPs and patients.

The next era of medical information is human-centric and agentic. By embedding AI directly into workflows, standardizing data structures, and balancing technology with human expertise, medical information can evolve into the trusted orchestrator of scientific knowledge in today's dynamic information ecosystem.

Take the next step toward agentic medical information and explore Agentic Case Processing in MedInquiry.



ABOUT US

Veeva is the global leader in cloud software for the life sciences industry. Committed to innovation, product excellence, and customer success, Veeva serves more than 1,100 customers, ranging from the world's largest biopharma companies to emerging biotechs. As a Public Benefit Corporation, Veeva is committed to balancing the interests of all stakeholders, including customers, employees, shareholders, and the industries it serves.

For more information, visit www.veeva.com

