GE HealthCare

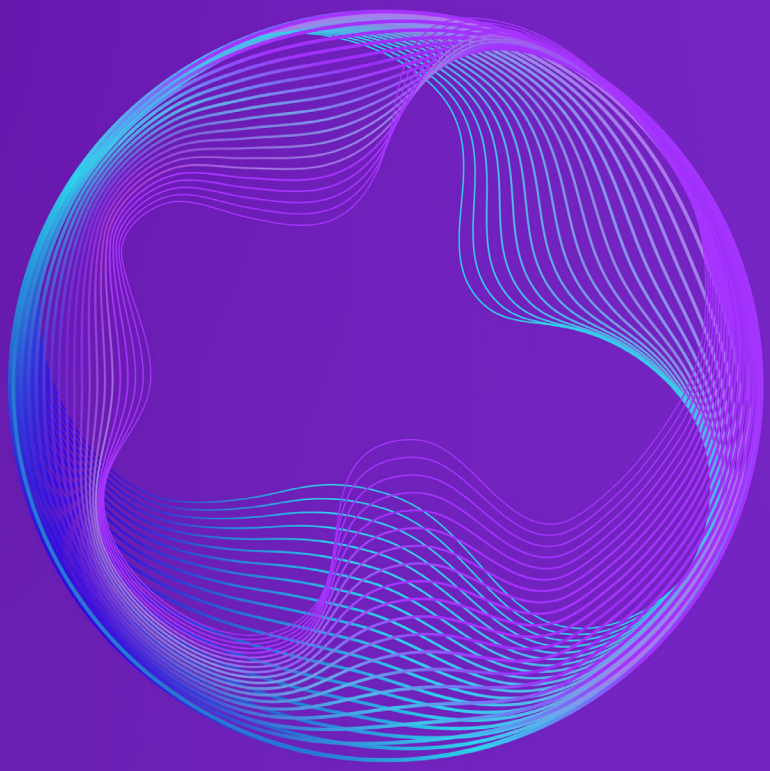

Agents of change: Exploring the next era of AI in healthcare

Healthcare is approaching an inflection point, where its greatest promise intersects with its most complex challenges. The industry generates enormous amounts of data—hospitals produce an average of 50 petabytes annually.[1] Yet an estimated 97% of this healthcare data goes unused, archived without being retrieved or analyzed.[2]

The human cost of this data overload is significant for both care providers and patients. Many clinicians now spend more time in front of electronic medical records than with patients. Healthcare organizations experience high staff turnover, and three in ten healthcare workers have reported considering leaving due to burnout. Staff shortages are increasing cognitive burden on providers, and the World Health Organization projects a shortfall of 11 million healthcare workers by 2030.[3] These pressures exist within a global context in which 4.5 billion people lack adequate access to essential care.[4]

Amid these challenges, advances in artificial intelligence offer a promising research direction. With the ability to analyze vast amounts of data, AI and foundation models may help healthcare systems identify inefficiencies and support clinicians in their daily decision-making.

Over the last decade, GE HealthCare scientists have used deep learning to make imaging devices more efficient. For example, AIR™ Recon DL for MR imaging has shown reductions of up to 50% in scan times while maintaining image quality.[5]

GE HealthCare's Command Center, a commercially available solution, uses predictive analytics to help improve hospital operations. At Deaconess Health System, it enabled treatment of about 2,000 additional patients annually without adding beds. Humber River Hospital reduced length of stay, creating an efficiency equivalent to 35 beds of capacity, through improved flow coordination.

Building on such experience, GE HealthCare researchers are now studying how new AI architectures could assist healthcare providers and operational teams through greater automation, context awareness, and adaptive reasoning.[6]

---

[1] https://www.weforum.org/stories/2024/01/how-to-harness-health-data-to-improve-patient-outcomes-wef24/#:~:text=Improving%20data%20capture%2C%20sharing%20and,the%20power%20of%20the%20analysis.
[2] https://www.beckershospitalreview.com/healthcare-information-technology/innovation/hospitals-only-use-3-of-data-microsoft-says/#:~:text=While%20hospitals%20collect%20massive%20amounts,goes%20unused%2C%20according%20to%20Microsoft.&text=The%20World%20Economic%20Forum%20estimates,care%20in%20a%20timely%20manner.%E2%80%9D
[3] https://www.who.int/health-topics/health-workforce#tab=tab_1
[4] https://www.who.int/news/item/18-09-2023-billions-left-behind-on-the-path-to-universal-health-coverage
[5] https://www.gehealthcare.com/-/media/gehc/us/files/products/magnetic-resonance-imaging/3-0t/case-study-air-recon-dl--jb19184xx.pdf?rev=-1
[6] Concept only. May never become a product. Not for Sale. Not cleared or approved by the U.S. FDA or any other global regulator for commercial availability.

The evolution toward agentic AI

Most existing AI applications in healthcare remain reactive, responding to clinician queries, flagging abnormalities, or accelerating existing workflows. Current research into “agentic” AI explores systems that could go beyond assistance to coordinate and adapt more proactively.

Agentic AI represents a new paradigm under exploration. Instead of waiting for commands, these systems are designed to perceive context continuously, analyze complex information, propose actions, and maintain memory across interactions. This evolution could eventually transform AI from a reactive tool into a more collaborative system, subject to rigorous safety, validation, and regulatory evaluation.

Five technical pillars under study

Agentic AI research focuses on five interconnected pillars that may enable capability, accountability, and safety.

01 Tool Use

Tool Use enables agents to interact with existing healthcare infrastructure, including electronic health records, laboratory information systems, imaging platforms, and decision-support tools. Experimental frameworks may use open standards, such as JSON Schema for universal tool description and Model Context Protocol for standardized interactions. These formats could reduce integration complexity and support interoperability across healthcare IT systems.

For example, in a simulated emergency scenario, an AI agent could query a hospital’s EMR for cardiac history, retrieve EKG results, access lab data, and review imaging from PACS—all within seconds. This concept demonstrates how structured protocols might allow for faster, more comprehensive data access across previously disconnected systems.

02 Multi-Hop

Multi-Hop Data Retrieval with Verification seeks to address the challenge of combining information from multiple sources of differing reliability. Agents under study simultaneously query clinical databases, research resources, and validated medical guidelines. Each data point is evaluated through confidence scoring based on source credibility and recency. Full provenance tracking allows researchers to trace how a system reached a conclusion or recommendation.

03 Memory Systems for Adaptive Intelligence

Memory Systems for Adaptive Intelligence are being explored to help agents maintain context across sessions. Traditional chatbots lose memory after each exchange, but these prototypes aim to maintain several types of memory: working memory for immediate context, long-term memory for accumulated domain knowledge, episodic memory for prior interactions and procedural memory for institutional workflows. Together, these components could allow AI systems to adapt to evolving information while preserving patient-specific context.

04 Dynamic Attribute-Based Access Control (ABAC)

Dynamic Attribute-Based Access Control (ABAC) is designed to ensure privacy and compliance. Each access request is evaluated using multiple criteria—user identity, role, patient relationship, purpose, and timing—to determine appropriate data visibility. This mechanism supports fine-grained access aligned with privacy regulations.

05 Data Sanitization and Immutable Audit Trails

Data Sanitization and Immutable Audit Trails provide transparency and traceability. All personally identifiable information is automatically removed before data processing. Each system action, query, or decision is recorded in a tamper-resistant log, supporting accountability and future audit reviews.

From one agent to many: Multi-agent collaboration

The potential of agentic AI may be realized through collaboration between multiple specialized agents, mirroring clinical teamwork. GE HealthCare's research project, Health Companion, is exploring how multiple AI agents might interact to support complex decision processes.[7]

Within this experimental architecture, a coordinating agent functions as a virtual attending, synthesizing input from domain-specific agents such as radiology, pathology, clinical data, or genomics. Each specialist agent contributes its expertise while sharing context through standardized protocols that preserve meaning and ensure consistent communication.

The system is being evaluated in research environments to determine whether such collaboration can replicate the structured reasoning and discussion of a multidisciplinary care team. Built with transparency and safety principles in mind, this work explores whether multi-agent systems could one day provide clinicians with more complete, explainable insights.

Agentic AI in action: Early research demonstrations at HLTH 2025

At HLTH 2025, GE HealthCare is demonstrating research prototypes exploring how agentic AI concepts could support hospital operations and clinical workflows.

Anticipating Hospital Operational Strain

Hospital operations are among the most complex logistical challenges in healthcare. With constrained resources and increasing patient demand, anticipating surges is critical.

GE HealthCare researchers are studying how multi-agent architectures could integrate with diverse hospital systems—including real-time EMR, bed management, scheduling, and imaging data—to identify pressure points before they occur.

Three prototype models are being investigated. A pressure forecast model uses the N-BEATS neural architecture to analyze historical trends and estimate potential operational strain up to 72 hours in advance. The model processes factors such as census data, emergency department activity, imaging turnaround times, and staffing levels to identify likely stress points.

An expected day of discharge (EDD) model applies fine-tuned BERT language models to estimate discharge timing. It processes structured data (labs, vitals) and unstructured data (clinical notes) to identify patient-level discharge patterns. A "sliding window" approach helps the model prioritize recent events while maintaining a full view of patient history.

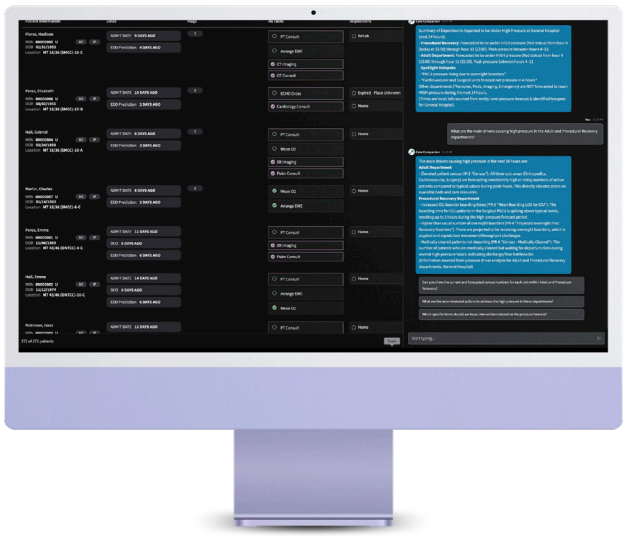

A conversational assistant prototype enables hospital leaders to ask operational questions in plain language and receive structured, traceable responses. This system is being studied to evaluate whether conversational AI can simplify access to operational insights and support decision-making through auditable, data-linked answers.

[7]Concept only. May never become a product. Not for Sale. Not cleared or approved by the U.S. FDA or any other global regulator for commercial availability.

Multi-agent AI for perinatal care: Supporting clinical decision research

Maternal and fetal outcomes remain a significant healthcare challenge. In the United States, severe maternal morbidity affects approximately 50,000 women annually, and maternal mortality rates remain among the highest in the developed world.[8]

GE HealthCare's research teams are conducting early-stage studies to explore whether AI could assist clinicians in managing large, multimodal data flows during labor and delivery.

One prototype investigates whether conversational AI could generate near-real-time patient summaries, highlight relevant protocols, and surface key guidance for obstetric teams. The system is being evaluated in simulated or controlled environments and is not used for clinical decision-making.

Four conceptual workflows are under examination:

1. Real-time patient summarization, which compiles multiple data streams into concise role-specific summaries.
2. Protocol-aware question answering, which converts static hospital protocols into searchable formats so clinicians can quickly find steps during emergencies.
3. Decision support for delivery timing, which coordinates reasoning across maternal and fetal factors to present structured trade-offs in simulated studies.
4. Multimodal data integration, which combines structured vitals, waveform monitoring, and clinical notes to produce unified data views.

These explorations aim to determine whether conversational AI could one day help clinicians synthesize complex information more efficiently—always under expert oversight and within defined research boundaries.[9]

Understanding and mitigating emerging risks

As autonomy in AI increases, understanding potential risks is essential. GE HealthCare research teams are investigating safeguards to ensure security, reliability, and compliance for agentic systems.

Areas of focus include authentication, identity verification, and protection against unauthorized system access. Emerging threats such as synthetic voice or text impersonation could challenge conventional authentication. Studies are exploring ways to strengthen identity validation and mitigate misuse.

Other research examines how to detect and manage unexpected system behaviors, sometimes known as emergent behaviors, which can occur in AI models under conflicting objectives. Investigators are evaluating containment strategies and boundary conditions to ensure systems remain predictable and transparent.

Novel cybersecurity risks are also under review. Agent-based systems introduce potential vulnerabilities such as prompt injection or data poisoning. GE HealthCare researchers are testing countermeasures in simulated environments to strengthen resilience.

A comprehensive defense-in-depth framework is being evaluated, combining:

- Monitoring and observability infrastructure, enabling detailed tracking of agent interactions and activity thresholds;
- Architectural integrity and boundary enforcement, using single-purpose agents and version-controlled schemas to limit permissions;
- Quality assurance and validation, applying rigorous pre-deployment testing and simulated API environments; and
- Model-level security controls, using classifiers to detect sensitive information, prevent generation of inappropriate content and monitor for adversarial inputs.

Together, these layers are designed to support robust system integrity and uphold healthcare privacy standards.[10]

---

[8] https://www.commonwealthfund.org/publications/issue-briefs/2021/oct/severe-maternal-morbidity-united-states-primer#:~:text=A%20richer%20understanding%20of%20maternal,avoided%20with%20timely%2C%20appropriate%20care.

[9,10] Concept only. May never become a product. Not for Sale. Not cleared or approved by the U.S. FDA or any other global regulator for commercial availability.

Conclusion: Responsible research for the future of healthcare

Agentic AI research represents an early step toward systems that could one day operate as intelligent collaborators within healthcare. Through the five design pillars—tool use, verified data retrieval, adaptive memory, dynamic access control, and immutable audit trails—researchers are studying how AI might navigate complex environments while maintaining accountability and safety.

Multi-agent prototypes mirror clinical collaboration, with specialized agents working together to synthesize data, anticipate challenges, and support clinicians. The comprehensive research into governance, cybersecurity, and explainability reflects GE HealthCare's commitment to advancing innovation responsibly.

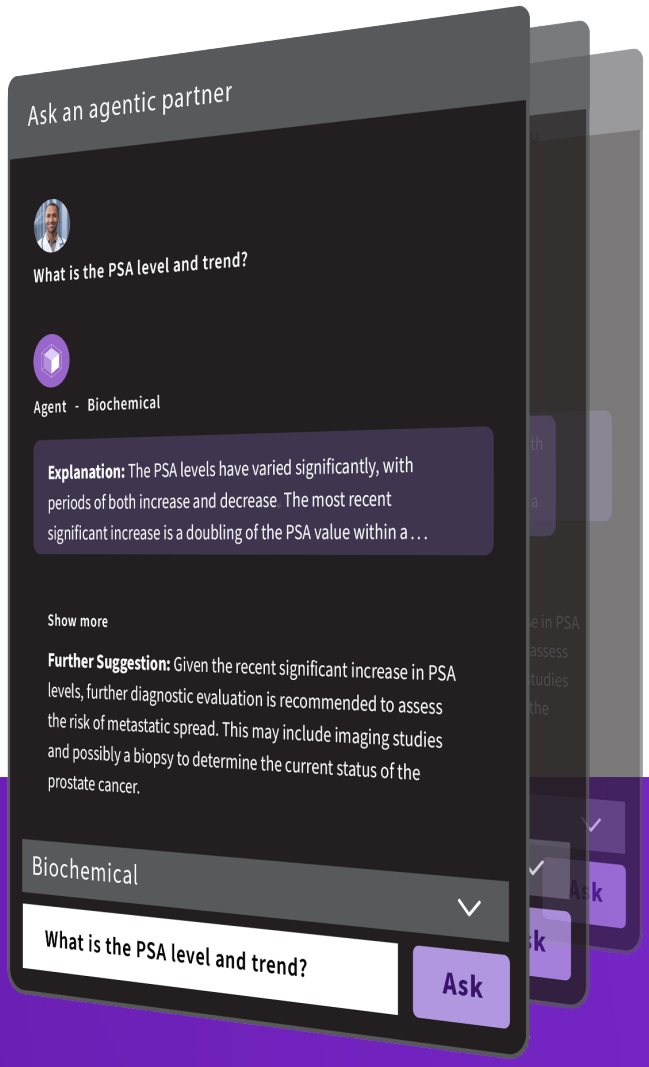

GE HealthCare

JB35391XX