The AI Catalyst Pulse

June 11th, 2025

Upcoming AI Catalyst Events

Mark your calendars

- **July 22: Mass Education: AI 101 for Healthcare Professionals** | Register here
    - Join a dynamic session of "AI 101 for Healthcare Professionals" designed to equip staff at member organizations with essential AI knowledge. Whether you're just starting to learn about AI or seeking to deepen your understanding, this session provides practical insights and guidance for leveraging AI effectively in healthcare settings.

Agentic AI in Healthcare: A No-Hype Executive Guide

AI vendors in healthcare suddenly can't stop talking about "agentic AI." Epic has announced an “agentic platform,” Microsoft is promoting a “healthcare agent service” in Copilot Studio, and Salesforce has launched an “AI agent platform” called Agentforce.

Still, when pressed for specifics, vendors offer wildly different definitions of what “agentic” really means — from basic workflow automation to autonomous decision-making systems.

Here's what you need to know to cut through the noise.

What exactly is ‘agentic AI,’ anyway?

“Agentic AI” refers to systems that can independently pursue goals across multiple steps, adapting their approach based on results. Unlike traditional AI that simply responds to prompts — such as when you ask ChatGPT a question — agentic systems can plan, execute tasks, and modify their strategies without constant human input.

This represents the latest evolution in AI capabilities. Just as machine learning moved beyond rules-based systems by learning patterns from data without explicit programming, agentic AI goes further by independently planning and adapting its approach to achieve goals.

In principle, AI “agents” could do any task that a human with access to a computer can do. But in healthcare, “agentic AI” most commonly refers, at least for now, to limited-purpose agents that help patients with well-defined, mostly administrative tasks, such as scheduling or preparing for appointments.

Are there different degrees or kinds of agentic AI?

Yes! The term “agentic” is immensely fuzzy. Think of it as being like “self-driving” technology in cars, a term that could refer to anything from basic lane-keeping assistance to a fully autonomous robotaxi.

Similarly, “agentic AI” capabilities range from basic workflow automation to fully independent decision-making. Here’s one useful, five-level framework for thinking about an AI tool’s degree of agency:

- **Level 1: AI-augmented automation.** These tools run through a predefined series of tasks, with some or all steps supported by AI.

- **Level 2: Agentic assistant.** These tools similarly run through predefined tasks, but they can call on other tools to provide enhanced capabilities, such as searching through files or browsing the internet.

- **Level 3: Planning and reflecting.** Rather than executing a fixed set of tasks, these tools can flexibly modify their task list, reasoning through unexpected obstacles to keep pursuing their end goals.

- **Level 4: Self-refinement.** These tools, which remain hypothetical, can actively improve themselves without human input, gaining new capabilities to flexibly pursue their goals.

- **Level 5: Autonomy.** These not-yet-developed tools represent artificial general intelligence (AGI), capable of essentially any task a human worker could complete with a computer – and potentially much more.

Current healthcare implementations cluster around levels 2 and 3. WellSpan's deployment of Hippocratic AI's Ana, for instance, demonstrates Level 3 autonomy: The system adapts conversations based on patient responses, handles English and Spanish speakers differently based on engagement patterns, and manages more than 98% of calls without human intervention.

Health systems are drawn to agentic AI because it can handle the kind of complex, back-and-forth problem-solving that previously required human staff. The ability to retrieve information, clarify needs, and iteratively work toward solutions makes agentic AI particularly valuable for patient communication and care coordination tasks that have historically consumed significant staff time.

Why does it matter whether an AI tool is 'agentic' or not?

For an AI agent to be useful, it must flexibly complete multiple tasks on its own. This means there's no "human in the loop" for routine decisions, removing what has historically been a key safety precaution for healthcare's AI implementations.

For low-risk, backend administrative tasks, this autonomy may not present any concerns. Even if an agent makes an error, it's unlikely to cause patient harm.

Patient-facing AI agents, however, raise bigger challenges. For example, MyEleanor at Montefiore Einstein conducts sensitive conversations about why patients missed cancer screenings. The system must navigate cultural barriers, health literacy challenges, and patient anxiety — all without human oversight for most interactions.

If you're exploring this latter kind of agentic AI tool, you'll need to ask new-in-kind questions: What decisions should these systems make independently? When must they escalate to humans? Do humans audit the agent's final choices?

Smart agentic implementations navigate these challenges by defining clear boundaries and metrics for success. For instance, MyEleanor escalates complex emotional situations to human staff, and it has achieved a 57% patient engagement rate and a 25% procedure completion rate.

Can AI agents ever collaborate with other AI agents?

This question highlights another key dimension in evaluating agentic AI: whether you're dealing with one AI agent or multiple agents working together.

- **Single-agent systems use one AI to handle an entire workflow.** With MyEleanor, for instance, one agent manages the complete patient outreach process. It initiates contact with patients who missed colonoscopy appointments, conducts a sophisticated conversation to identify barriers (assessing 14 different potential obstacles from transportation to work conflicts), helps problem-solve those barriers, facilitates rescheduling, and documents the interaction.

- **Multi-agent systems coordinate multiple specialized AIs, each handling specific subtasks.** Think of this as a relay race rather than a solo sprint. An example is Notable's platform, which uses different agents to handle contact center operations, clinical documentation, and care coordination. They work together across the revenue cycle, from patient intake through claims processing.

Single agents work well for contained workflows with clear start and end points, such as patient outreach or document processing. Multi-agent systems are more appropriate when workflows cross departmental boundaries or require diverse expertise. They're harder to implement initially, and may pose even more challenging governance and safety questions, but easier to build on over time.

How should I evaluate vendors who are offering agentic AI solutions?

We've previously shared our FUMBLE Framework for cutting through AI vendor hype, as well as UVM's questionnaire for vetting AI vendors. In addition to using those tools to assess agentic AI vendors, consider asking these questions:

- **Autonomy and decision-making:**
    - "What actions can your system take without human approval?"
    - "Can you walk me through a specific example of how your system handles an exception?"
    - "What percentage of cases require human intervention, and why?"
    - "How can we audit decisions made autonomously by your system?"
- **Architecture and integration:**
    - "Draw your system architecture — is this one agent or multiple?"
    - "What happens when an agent fails or produces an unexpected result?"
    - "Which existing systems does this tool need to access, and what actions can it take within those systems?"
- **Governance and oversight:**

- "What monitoring capabilities come built-in?"
- "How do we set and modify the boundaries of agent autonomy?"
- "What's your recommended governance structure for this level of autonomy?"
- "How do other clients handle liability for autonomous decisions?"

What strategic questions should our executive team be considering about agentic AI?

Here's a list of prompts to get you started:

- **About governance readiness:** "Do our current AI oversight structures work when humans aren't reviewing every decision?" Most organizations need to reevaluate their oversight processes once AI reaches Level 3 autonomy and begin independently structuring their own decision-making processes.

- **About liability models:** "Who's accountable when an autonomous agent makes a mistake?" Current malpractice frameworks assume human decision-makers. Legal and risk teams need involvement before, not after, implementation.

- **About success metrics:** "How do we measure success beyond efficiency?" MyEleanor improved patient compliance, not just call volumes. Ana revealed unmet demand in Spanish-speaking populations. The most valuable outcomes may be unexpected.

- **About workforce impact:** "How should we message our use of agentic AI to staff who may fear displacement?" WellSpan succeeded by involving nurses in selecting use cases, positioning Ana as handling overflow work rather than replacing existing roles.

- **About timing:** "When will agentic AI become table stakes?" Early adopters will face integration challenges but stand to gain competitive advantages. Later adopters will get more mature products but may struggle to differentiate themselves from competitors or may lag in offering their patients much-needed capabilities.

What's the bottom line?

Agentic AI matters not (just) because it's more powerful than previous AI, but because it fundamentally changes *who* makes decisions. For the first time, AI systems can handle entire workflows independently, from initial patient contact through problem resolution.

This shift requires you to ask harder questions than demanded by past AI implementations. Not just "Is it accurate?" but "Should it operate autonomously?" Not just "Does it integrate?" but "Who's accountable when it acts independently?"

The health systems that succeed with agentic AI will be those that meet these powerful new technologies with a similarly sophisticated degree of organizational readiness.

Are you implementing and using agentic AI tools at your organization? Or has your team developed a successful questionnaire to vet agentic AI solutions? Let us know! We'd love to hear about your experiences at aicatalyst@hmacademy.com.

Is AI Really Taking Entry-Level Jobs? What Healthcare Leaders Need to Know

The New York Times, *The Atlantic*, and *Axios* are all telling the same story: AI is decimating entry-level white-collar jobs. Unemployment among recent college graduates has risen to 5.8%, prompting the New York Fed to warn that labor conditions for young workers have "deteriorated noticeably." Anthropic's CEO, Dario Amodei, went further, predicting AI could eliminate half of all entry-level white-collar jobs within five years.

So is this entry-level job apocalypse real? And is it about to wallop the healthcare workforce?

What the data shows so far

We'd argue there's a case for skepticism:

- **Yes, data shows an uptick in entry-level unemployment … but it's a small jump in a noisy series.** The recent graduate unemployment has risen by about one percentage point since last year. Still, its current level is less than half the peak hit during the Covid-19 crisis, and as *The Atlantic* notes, the upward trend could also reflect post-pandemic hiring freezes or broader economic challenges.

  Take a look at the chart below. This job-loss narrative is driven by the tiny uptick on the far right in the light blue line. It's real, but probably not worth panicking over yet:

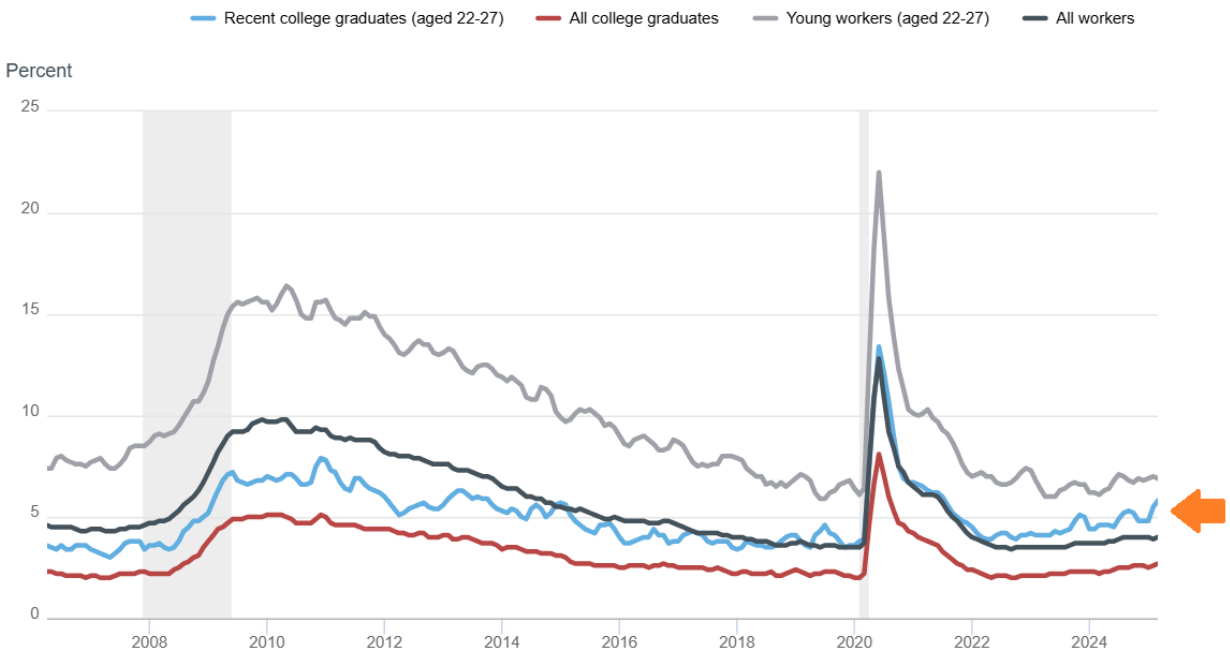

- **Entry-level healthcare work is particularly automation resistant.** In 2023, OpenAI released a study scoring 19,000 workplace tasks on their vulnerability to AI automation. Unsurprisingly, their data suggested that hands-on jobs and highly regulated jobs – like many in healthcare – would

be challenging to automate. This would intuitively seem all the more true of entry-level healthcare jobs, which so often involve moving patients, operating equipment, and delivering hands-on care.

To test this intuition, we applied OpenAI's methodology to healthcare-specific roles. Their original study scored individual work tasks from 0% to 100% automatable; we aggregated these scores by occupation, weighting core job functions more heavily than supplemental ones.

We found that only 29% of entry-level healthcare tasks (which we defined as those requiring no on-the-job experience and no more than a bachelor's degree) are vulnerable to AI disruption, compared to 35% of more experienced healthcare work. By contrast, 56% of entry-level legal work is vulnerable.

Still, some healthcare jobs really are vulnerable — and we already have proof

You might doubt whether OpenAI's study can really predict future employment trends – especially given that it is, by now, two years old. But trends since the study's publication have borne out its predictions.

By our analysis of OpenAI's data, the single most AI-vulnerable job in healthcare was "medical transcriptionist." And sure enough, today's most widely deployed healthcare AI tools are clinical documentation assistants such as Nuance/DAX and Abridge.

So what other entry-level healthcare jobs look especially vulnerable by this methodology? Here are the top five:

- Medical records specialists
- Diagnostic medical sonographers
- Health information technologists and medical registrars
- Medical assistants
- Pharmacy aides

Notably, many of these fields currently face acute *shortages* of human workers. As AI increasingly handles similar tasks, the immediate effect might not be humans losing jobs but, rather, humans becoming more efficient.

One caveat: Our analysis of "healthcare work" focuses only on jobs formally categorized as healthcare-related in federal databases. Many entry-level jobs at health systems fall into other labor categories. For instance, medical secretaries, billing clerks, and insurance processors are all "office and administrative" jobs — and all score above 50% automatable.

If an entry-level job crisis really is gaining speed, it might hit these back-office operations first.

Questions to consider:

1. Which regulatory, workflow, or reimbursement barriers will slow automation of technically automatable healthcare tasks? Should you work to preserve or remove these barriers?

2. How should your AI deployment strategy differ between clinical roles and administrative roles?

3. Looking at the high-automation list above, which roles would benefit from AI use to address shortages, and which might face genuine pressure to displace human workers?

Our Four Critical Steps to Transform Capacity Management with AI

Over 90% of hospital leaders cite staffing constraints as their top challenge. Nearly 84% point to inefficient discharge processes. Meanwhile, patients wait in ED hallways while OR blocks sit empty and discharge-ready patients occupy acute beds for days.

AI-powered capacity management promises to break these logjams through predictive analytics, automated workflows, and real-time optimization. But results vary dramatically. Some health systems report millions in savings within months, while others struggle to achieve meaningful results.

Our latest briefing, "Four Critical Steps to Transform Capacity Management with AI," identifies what separates breakthrough results from marginal improvements. Here are the four key strategies we identified:

1. **Fix your processes before you invest in AI — otherwise, you'll just automate a bad practice.** OhioHealth assessed six hospitals for Qventus implementation but deployed at only two. Those sites had standardized discharge rounding processes; the others showed too much workflow variation. Scott Estep, SVP of Nursing Operations, explains: "AI is great, but if you don't build the process to support it, the technology is just noise." The two implementation sites reduced excess days by 20% and saved $1.7 million and recouped investment costs within six months.

2. **Hardwire cultural adoption into workflows; staff won't use AI that feels optional.** Nebraska Medicine's discharge lounge handled just five patients weekly despite available capacity. But after they embedded AI recommendations directly into required nursing workflows – making the system automatically identify eligible patients and integrate with existing processes – utilization jumped 2,000%. Discharge times dropped by 60 minutes system-wide.

3. **Use AI to prevent future bottlenecks, not just to report on past problems.** Geisinger's prior OR scheduling tools were leading to underutilized block time, staffing inefficiencies, and scheduling misalignments. To solve the problem, they implemented an AI-powered solution, Opmed.ai, to predict conflicts before they occur. The system now outperforms EHR predictions by 30%, accurately forecasts case durations 96% of the time, and frees 185 OR hours monthly through proactive optimization.

4. **If in doubt, target your AI investments at primary care and discharge planning first.** Presbyterian Healthcare Services piloted RhythmX AI in primary care, where high patient volumes multiply efficiency gains. The platform surfaces evidence-based recommendations during visits, enabling providers to address multiple conditions efficiently. Early data from RhythmX AI-executed research suggests providers can accommodate 3-4 additional appointments weekly while capturing $1-2 million in previously missed specialist referral revenue.

The health systems achieving transformative results aren't necessarily those with the most sophisticated technology. They're those that understood AI as a tool that reveals hidden connections across their entire system – but only when implemented thoughtfully within ready organizations.

Read our complete research brief for more details on these four key strategies, archetypes of organizations implementing these AI tools, and case studies from successful implementations.

AI Strategy Quick Hits

Noteworthy moves from peers to implement AI technologies

- **Mercy** responsibly advances and scales AI across clinical and operational domains
- **Privia Health** sees major gains from AI-enabled EHR tools from athenahealth and Navina
- **Mass General Brigham** expands ambient documentation to over 2,700 users
- **Northwell Health** tackles supply chain waste with **AssistIQ**
- **Cleveland Clinic**, **Mercy**, and **Community Health** modernize data management and physicians credentialing with AI
- **Intermountain Health** teams with **Perimeter** on AI breast cancer imaging
- **Griffin Health** improves cancer screening follow-up with **Inflo Health's** AI-powered platform

Emerging Use Cases

New capabilities that indicate AI's potential

- **Mass General Brigham** pairs its diagnostic decision support system with LLM's to improve diagnostic accuracy
- **St. Luke's** and **LVHN** deploy diverse AI tools for cancer and cardiac care
- New AI test helps predict which prostate cancer patients benefit from abiraterone

Market Moves

A round-up of AI company announcements and stories

- **Epic** launches Launchpad to speed generative AI deployment
- **Microsoft** launches an AI agent orchestrator within its Azure AI Foundry platform designed specifically for cancer care coordination
- **St. Luke's University Health Network** invests $30M+ in **GE HealthCare**'s AI MRI platform
- Innovaccer launches its Gravity platform that unifies siloed data for value-based care (VBC) initiatives
- **The Coalition for Health AI (CHAI)** establishes its first certified partnership for AI model validation

Policy Updates

Understanding the evolving AI regulatory and legislative landscape

- Thomas Keane, MD, has been appointed as the new national coordinator for health IT at HHS
- The **FDA** begins agency-wide deployment of an AI tool, Elsa, intended to expedite clinical protocol reviews and scientific regulatory evaluations

Expert Insights

For further reading, articles, videos, and podcasts that we found insightful

- *Becker's Health IT*, "Health systems add, drop roles with AI"
- *Psychology Today,* "Enhanced LLM Aces the US Medical Licensing Examination"