AI on Friday

What to Know, Why It Matters

Issue 11, 6 November 2026

Pan American Health Organization
World Health Organization Americas Region

Information Systems and Digital Health Unit
Department of Evidence and Intelligence for Action in Health

This weekly email shares key developments, use cases and practical AI-in-health concepts in plain English, helping us apply AI in our technical cooperation and discuss emerging opportunities with Member States and partners.

DEVELOPMENT TO WATCH

A $62.7 million effort to develop an FDA-authorized clinical AI agent

What is happening: The U.S. Advanced Research Projects Agency for Health has selected teams for its four-year ADVOCATE program. The initiative will develop patient-facing AI agents for cardiovascular care, supervisory agents to monitor their behavior and a plan for deployment across different health settings.

Potential value: The agents are expected to provide continuous support with appointments, medications, diet, exercise and other aspects of cardiovascular care. They could extend specialist capacity, particularly in rural and underserved communities.

What to keep in mind: These systems are still under development and have not received FDA authorization. Allowing an AI agent to recommend or initiate clinical actions requires rigorous validation, secure access to health data, clear escalation to professionals, continuous monitoring and defined accountability.

Why it matters: ADVOCATE moves the discussion from AI that provides information to AI that may take or coordinate clinical actions. Its use of supervisory agents also introduces an important model: one AI system monitoring another for unsafe recommendations or unexpected behavior.

Read more here

AI IN PLAIN ENGLISH

AI Agents

AI agents are systems that can pursue a goal through a sequence of steps. They can interpret a request, develop a plan, use digital tools or data, take actions, review the results and adjust what they do next. A chatbot usually responds to one prompt, while an agent can continue working until a task is completed.

Example: An agent receives the goal of organizing a meeting, checks participants' availability, proposes a time, prepares an agenda and sends invitations after receiving approval.

AI Agents in Health

AI agents in health can support multistep clinical, administrative and public health workflows. They may gather information from authorized sources, identify actions, prepare recommendations, coordinate tasks and monitor whether follow-up occurs.

Why it matters for health: These systems could reduce repetitive work and help prevent important tasks from being missed. Their actions may also affect patients, professionals and health services. They therefore require a clearly defined purpose, validated performance, secure access to data, human authorization, audit trails and mechanisms to pause or escalate decisions.

Example: An agent identifies patients overdue for follow-up, prepares personalized outreach messages and proposes appointments for staff approval.

Contributions *to this issue were received from Dr. Maka Tsulukidze of Florida Gulf Coast University, Associate Professor | Health Services Researcher | AI Strategy & Leadership in Higher Education & Healthcare, Florida Gulf Coast University · University of North Carolina at Charlotte.*

Important reminder: AI-generated content may contain errors, outdated information or incorrect references. Always verify original sources, publication dates and links before using or sharing it, particularly for clinical, policy or technical decisions.

FROM THE AMERICAS

Canada - Guide on the Use of Agentic Artificial Intelligence

In October 2026, the Government of Canada published **guidance for federal departments and agencies considering the use of agentic AI**. The guide explains that agents can sequence tasks, interact with digital systems and take actions, creating additional risks beyond those associated with generative AI. It calls for clear limits, human oversight, auditability and attention to privacy, security, fairness, human rights and accountability. These safeguards are especially relevant when agents are used in health and other sensitive public services. **Full text** here

Four Levels of Agent Autonomy

The Canadian guide describes a progression from limited assistance to greater autonomy:

Level 1, Assistive: The AI suggests options, while a person makes the decision and takes the action.

Level 2, Semi-autonomous: The agent proposes an action or prepares it for human approval.

Level 3, High autonomy: The agent completes delegated tasks within defined permissions, records its actions and informs the user.

Level 4, Adaptive autonomy: The agent monitors changing circumstances, adjusts its actions within established limits and escalates exceptions.

Before, During and After Use

Before use: Define a narrow purpose, limit the data, tools and permissions available to the agent, begin with read-only access and test realistic and adversarial scenarios.

During use: Assign a human owner, maintain logs of actions and approvals, monitor errors and automation drift, and keep an external mechanism to pause or deactivate the agent.

After deployment: Evaluate performance, bias, user feedback and impact. Reassess the safeguards whenever the agent's data, tools, permissions or purpose change. When an agent is retired, remove its access, preserve required records and document the lessons learned.

Additional recommended readings

Towards Autonomous Medical Artificial Intelligence Agents

Published in Nature in 2026, this study evaluated MIRA, an autonomous agent operating in a simulated electronic health record. Across more than 500 emergency-care cases, the agent gathered histories, ordered and interpreted tests, developed diagnoses and proposed treatments, procedures and admissions. Results were promising, but prospective evaluation in real clinical settings is still required. ***Read more*** *here*.

The Role of Agentic Artificial Intelligence in Healthcare: A Scoping Review

Published in npj Digital Medicine in 2026, this review identified only seven eligible studies across emergency medicine, oncology, radiology and rehabilitation. Most were exploratory and only one involved patients. The findings show the promise of AI agents while emphasizing the need for clearer definitions, standardized evaluation, clinical validation, regulatory guidance and professional oversight. ***Read more*** *here*.

Important reminder: AI-generated content may contain errors, outdated information or incorrect references. Always verify original sources, publication dates and links before using or sharing it, particularly for clinical, policy or technical decisions.