AI MedAgent

Continuous AI medication oversight for acute inpatient care

> **A demonstration of a patent-pending method. Advisory only. Not a medical device.**

What it is

AI MedAgent is a working demonstration of a patented method for continuous, predictive medication-safety reasoning in acute inpatient care. Its hypothesis is simple to state: that optimizing ongoing therapy proactively, rather than waiting to react, can benefit both the patient and the care team, and that a large language model reasoning over the live medication context, with external public data marts folded into that reasoning, is a credible way to do it. Better clinical workflow, tighter medication management, and improved patient outcomes are real potential gains from that approach. None of them is claimed or measured here. The scope of this demonstration is narrower and deliberately honest: to show that the method itself is doable and credible. The patient shown is simulated. The reasoning is real, generated live by Anthropic's Claude over the exact values on screen and current federal drug-safety data, and every cycle ends at physician review.

The method, which is the protected asset

The claim is not a device and not a dataset. It is a method: real-time reasoning by a large language model over the evolving medication context of one patient, on the continuous stream, drawing on public federal drug-safety data to catch a developing problem earlier than a threshold alarm or a scheduled review would. A reactive workflow waits for an absolute number to be crossed. This watches the trajectory. Proving that the approach improves outcomes would require proper blinded study; that remains a testable hypothesis, not a settled result, and saying so plainly is part of the design.

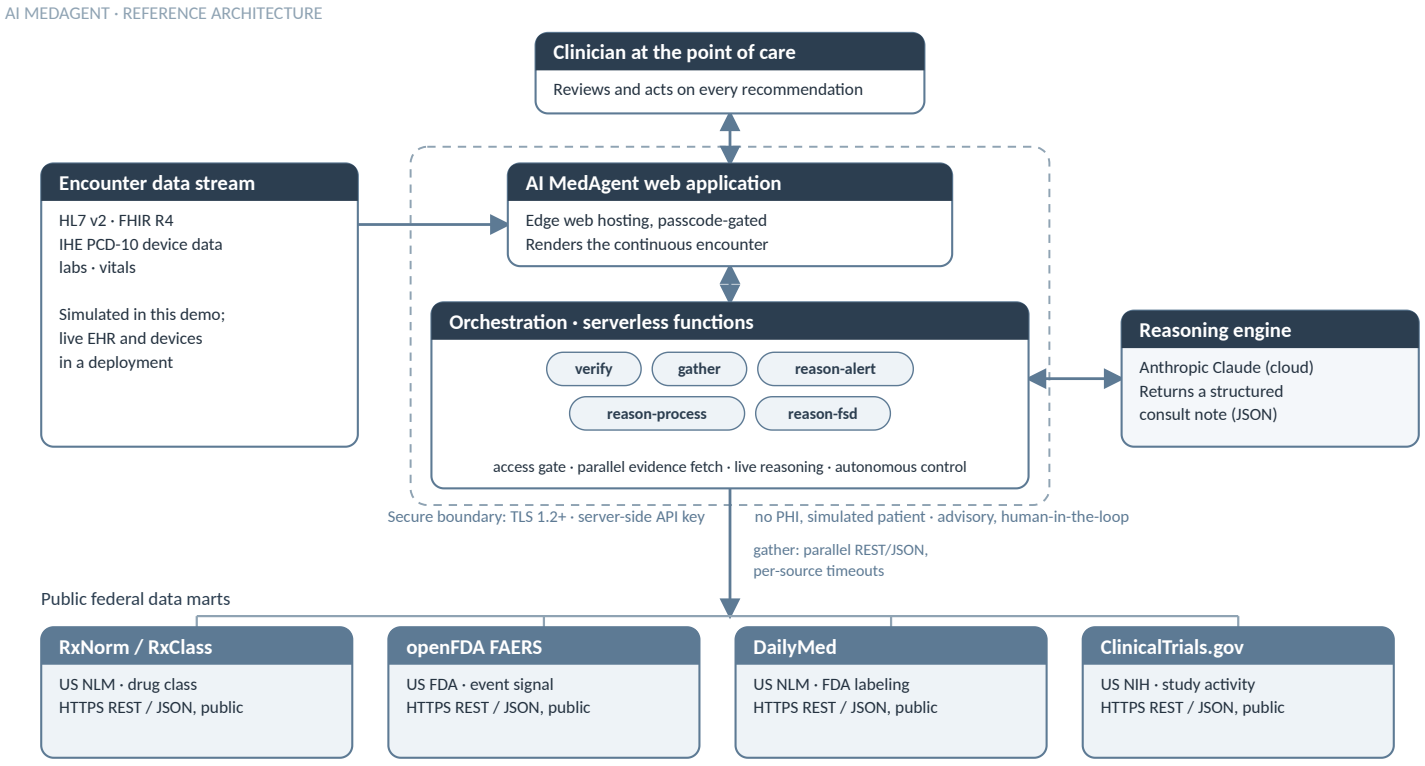

AI MedAgent reference architecture. Encounter data is simulated in this demonstration; a deployment would draw on the live EHR and connected devices. Reasoning is performed by Anthropic's Claude in the cloud, over evidence gathered from four public federal data marts.

How the data flows

1. Encounter in. A clinical stream (HL7 v2, FHIR R4, IHE PCD-10 device data, labs and vitals) drives the encounter. Simulated in this demonstration; the live EHR and connected devices in a deployment.

2. Evidence gathered once. On start, the gather function queries four public federal data marts in parallel, each with its own timeout, and caches the result for the encounter. One slow or unavailable source cannot stall the loop.

3. Reasoning, live. At each moment of interest, the reasoning functions send the exact patient values plus the gathered evidence to Anthropic's Claude, which returns a structured consult note: the risk, the reactive-versus-predictive contrast, the traceable reasons, the supporting evidence, options for the clinician, and an honest statement of what it could not see.

4. Physician review closes the loop. In advisory mode, nothing is administered or ordered. Every recommendation is presented for a clinician to accept, adjust, or reject.

5. A separate autonomous mode, documented separately. A fifth function, reason-fsd, is the controller for Full Self Delivery, a supervised autonomous mode in which the same reasoning decides and acts within clinician-ordered bounds rather than only advising. It is never engaged in the advisory case described here. The scope, the safety architecture, and the honest limits of that mode are set out in the FSD addendum to this datasheet.

Public federal data sources

SOURCE	PROVIDER	WHAT IT CONTRIBUTES	ACCESS
RxNorm / RxClass	US National Library of Medicine	Drug identity and drug class	HTTPS REST / JSON, public
openFDA FAERS	US Food and Drug Administration	Population adverse-event signal	HTTPS REST / JSON, public
DailyMed	US National Library of Medicine	FDA structured product labeling	HTTPS REST / JSON, public
ClinicalTrials.gov	US National Institutes of Health	Research-activity signal, cited as activity, not outcome	HTTPS REST / JSON, public

Security and governance

- Access is passcode-gated; the passcode is checked server-side and never sits in the page.
- The Anthropic API key is held server-side and is never exposed to the browser.
- All traffic is TLS 1.2 or higher.
- No protected health information. The demonstration patient is synthetic.
- Advisory only, human-in-the-loop. In the advisory mode described here, the method never administers or orders. The supervised autonomous mode is documented separately in the FSD addendum.
- The demonstration runs under a fixed monthly budget. When the cap is reached or a call cannot complete, a validated fallback note is shown and labeled, so the display is always honest about whether a response was live.

What it is not

AI MedAgent is not a medical device, does not diagnose, and is not a production drug-diversion surveillance system, which would train on millions of transactions. It is a demonstration of a method, built to be evaluated on its merits.

Patent and provenance

Patent pending. The method is the subject of a US provisional patent application filed in 2026. It continues one long thread of the same idea: real-time reasoning over a live physiological signal (a processed-EEG algorithm, 1986), population-scale medical-terminology reasoning (UMLS on the MediCompass platform, 2003), and now language-model reasoning over the live medication context.

For additional information, contact Daniel Pettus at danpettus@outlook.com.