

# Developing and Deploying AI in Healthcare



# Presenters



**Sean  
McGunigal**

Cognitive Computing Developer

# Agenda

**01** Why is AI in healthcare challenging?

**02** What can you do to be successful?

**03** How the landscape of AI is changing.



# Setting the Stage

November 2022

ChatGPT released



April

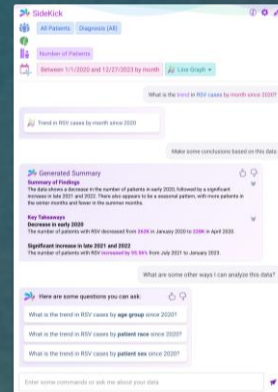
First customers live with In Basket Art

HIMSS23

Epic integrates generative AI into EHRs

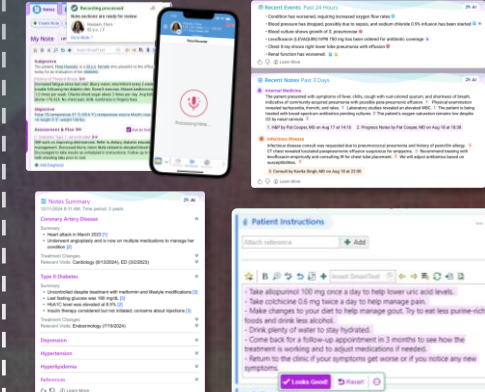
August

SideKick agent live on Cosmo



Summer

IP Insights, AI Text Assistant, Sidekick, Discharge Summary, Ambient Flowsheets, End of Shift, PB Coding Assistant live



Today

415 Epic customers using generative AI

Summer

150+ use cases live & in development

Winter

Cogito Cloud live  
AI Data Foundation

nature

May

Epic supports GPT-4

March

Epic integrates LLMs

Spring

Analytics dashboard insights & AI-extracted radiology follow-ups live

1<sup>st</sup> international customers live in Dutch

August

Live demo of MyChart Patient Experience Agent at Epic Users Group Meeting

Spring

Launchpad adoption program



# Agenda

**01** Why is AI in healthcare challenging?

**02** What can you do to be successful?

**03** How the landscape of AI is changing



# Why is AI in healthcare challenging?

## For Model Developers

Healthcare data is disparate and challenging to model on

Data can vary widely between organizations

PHI/PII makes it hard to do real-time data capture for error analysis

Test and re-test (even small changes can have large implications)

# Why is AI in healthcare challenging?

## For Model Developers

Healthcare data is disparate and challenging to model on

Data can vary widely between organizations

PHI/PII makes it hard to do real-time data capture for error analysis

Test and re-test (even small changes can have large implications)

## For Healthcare Organizations

AI requires iteration, speed, and agility; not all organizations are suited for that.

Successful groups are thoughtful about *post-release*. Who's monitoring? How?

Workflows can vary between hospitals; differentiate between on-label vs off-label.



# Why is AI in healthcare challenging?

## For Model Developers

Healthcare data is disparate and challenging to model on

Data can vary widely between organizations

PHI/PII makes it hard to do real-time data capture for error analysis

Test and re-test (even small changes can have large implications)

## For Healthcare Organizations

AI requires iteration, speed, and agility; not all organizations are suited for that.

Successful groups are thoughtful about *post-release*. Who's monitoring? How?

Workflows can vary between hospitals; differentiate between on-label vs off-label.

## What counts as a medical device?

You need to be aware of the regulations (which continue to change)

The more the AI tool does on its own, the more likely it will be a medical device

Risk primarily comes from the use-case, not necessarily the type of model



# Agenda

**01** Why is AI in healthcare challenging?

**02** What can you do to be successful?

**03** How the landscape of AI is changing

# Understand the *Full* User-Story

- **Focus on the workflow**
  - What data is available?
  - When is the user “ready” for the information?
  - What should be shown?
- **Supporting information helps build trust**
- **Trust leads to adoption**

## Notes Summary



Generated at: 6/3/2024 2:50 PM. Focused on: CHF

- Hussain, Flora, 52, has a history of congestive heart failure due to hypertension, which is currently stable. [1](#)
- She is on a treatment regimen that includes dietary sodium restriction, regular aerobic exercise, and daily weight measurement. [2](#)
- Her heart failure is to be reassessed in 3 months. [2](#)
- She is also managing diabetes, which is stable and will also be reassessed in 3 months. [1](#)
- Her current medications are well-tolerated without significant side effects. [3](#)
- She is exercising 5 days a week and is compliant with her hypertension medications. [2](#)
- A pre-visit phone call was completed on 5/22/2024, during which she reported moderate activity. [4](#)

## References




   Learn More



# Start Small, Build Out



Don't boil the ocean;  
a small solution can  
have a big impact on  
users.



Small use-cases lead to  
clear requirements and  
goals.

This makes it easier to  
iterate.



Pilot users → broad  
rollout

- Specialty, focus area
- User role
- Super users

# Build, Measure, Iterate

Don't focus on perfection out-of-the-gate, but do **evaluate at every stage**.

How will you measure your tool and **what are your "success" criteria?**

How will you **work with your customers & data partners** (who has the data/validation) on an **ongoing basis** to validate?

**Tight develop-test loops** will help you move faster.



# Agenda

**01** Why is AI in healthcare challenging?

**02** What can you do to be successful?

**03** How the landscape of AI is changing

# The Evolution of AI




## Rule-Based Logic

*Expert-defined and explicitly coded*



## Predictive Analytics


*Statistically derived to predict a pre-defined event*



## Generative AI

(Large Language Models)

*Generally-trained to generate novel content*



**Deterministic**

Targeted

**Probabilistic**

More generalized




# The Evolution *of* Generative AI



Single-Step




Multi-Step




Reasoning Agent



Dynamic Agent



The length of tasks AI can do is **doubling every 7 months.**





# Why Language Models Hallucinate

Adam Tauman Kalai\*  
OpenAI

Ofir Nachum  
OpenAI

Santosh S. Vempala<sup>†</sup>  
Georgia Tech

Edwin Zhang  
OpenAI

September 4, 2025

## Abstract

Like students facing hard exam questions, large language models sometimes guess when uncertain, producing plausible yet incorrect statements instead of admitting uncertainty. Such “hallucinations” persist even in state-of-the-art systems and undermine trust. We argue that language models hallucinate because the training and evaluation procedures reward guessing over acknowledging uncertainty, and we analyze the statistical causes of hallucinations in the modern training pipeline. Hallucinations need not be mysterious—they originate simply as errors in binary classification. If incorrect statements cannot be distinguished from facts, then hallucinations in pretrained language models will arise through natural statistical pressures. We then argue that hallucinations persist due to the way most evaluations are graded—language models are optimized to be good test-takers, and guessing when uncertain improves test performance. This “epidemic” of penalizing uncertain responses can only be addressed through a socio-technical mitigation: modifying the scoring of existing benchmarks that are misaligned but dominate leaderboards, rather than introducing additional hallucination evaluations. This change may steer the field toward more trustworthy AI systems.

# Epic AI Trust and Assurance Suite



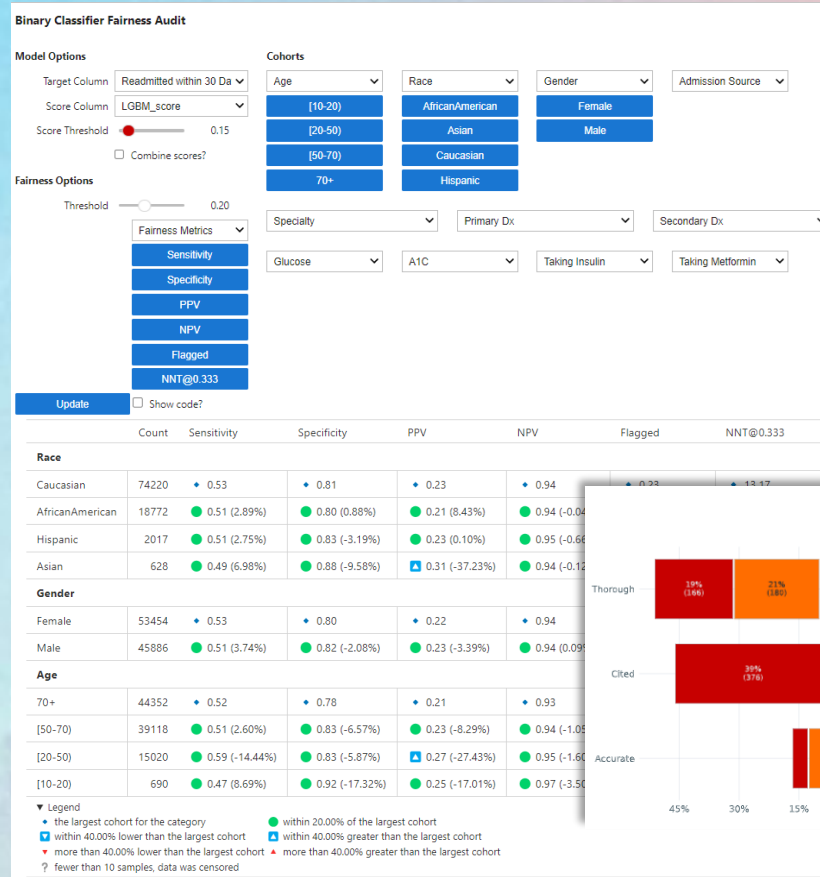
**Open-source** template and schema




Performance on **local patient mix**




Real-time, **ongoing monitoring**



check it out



Thank you to the following customers for contributing!



Health Care

UC San Diego Health




# CDS HOOKS™



AVAILABLE NOW

***Ep*ic**



Clinical Decision Support  
(CDS) Service



OurPractice Advisory - Bala, Chandra

**⚠ This patient has exceeded their lifetime dose of radiation exposure.**  
[Click here for ACR Radiation Guidelines](#)

**⚙ Perform the following actions**

Remove the following orders

☒ Remove: 🏠 CT SCAN HEAD CONTRAST  
Routine, Normal

Apply the following

☒ Order: 🏠 MRI Head w and wo IV Contrast

☐ **▶ Acknowledge and continue**

✓ **A**cept (2)

✗ **C**ancel

# Generalize *to* Benefit More *with* Industry Standards

## CDS Hooks



<https://cds-hooks.hl7.org/>

**CDS Hooks**  
2.0.1 - STU 2 Release 2

[CDS Hooks](#) [Hook Library](#) [Change Log](#) [Downloads](#)

[Table of Contents](#) > [Home](#) <prev | bottom | next>

This page is part of the CDS Hooks Specification (v2.0.1: STU2) based on FHIR (HL7® FHIR® Standard) R4. This is the current published version. For a full list of available versions, see the [Directory of published versions](#)

## 1 Home

Official URL: <a href="http://cds-hooks.hl7.org/ImplementationGuide/hl7.fhir.uv.cds-hooks">http://cds-hooks.hl7.org/ImplementationGuide/hl7.fhir.uv.cds-hooks</a>	Version: 2.0.1
IG Standards status: <a href="#">Trial-use</a>	Computable Name: CDSHooks

### 1.1 Intellectual Property Statements

The HL7 CDS Hooks Implementation Guide is the copyright of HL7 International and Boston Children's Hospital. The specification is licensed under a Creative Commons Attribution 4.0 International License.

### 1.2 Overview

This HL7 CDS Hooks Implementation Guide is published at the level of [Standard for Trial Use](#). It describes a "hook"-based pattern for invoking decision support from within a clinician's workflow. The API supports:

- Synchronous, workflow-triggered CDS calls returning information and suggestions
- Launching a user-facing SMART app when CDS requires additional interaction

The companion [HL7 CDS Hooks Library](#) contains specifications of industry standardized clinical workflow steps used by systems conforming to this guide. While changes to this guide become infrequent and tightly constrained, new hooks will continue to be specified and matured in the Library.

See <https://cds-hooks.org/> for additional information, resources and ways to get involved.

### 1.3 Conformance Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this specification are to be interpreted as described in [RFC2119](#). Further, the key word "CONDITIONAL" indicates that a particular item is either REQUIRED or

- [Intellectual Property Statements](#)
- [Overview](#)
- [Conformance Language](#)
- [Use of JSON](#)
- [CDS Hooks Anatomy](#)
  - [CDS Services](#)
  - [CDS Clients](#)
  - [Cards](#)
- [Discovery](#)
  - [HTTP Request](#)
  - [Response](#)
  - [HTTP Status Codes](#)
  - [Discovery Example](#)
- [Calling a CDS Service](#)
  - [HTTP Request](#)
    - [hookInstance](#)
  - [Example](#)
- [Providing FHIR Resources to a CDS Service](#)
  - [Prefetch Template](#)
    - [Prefetch tokens](#)



cogito  Clarity

ergo sum

Kit 

Run analytics and  
extracts on

Caboodle

# Wrap Up

- Understand the full user-story and have a plan for all stages of the lifecycle (dev, test, release, post-live, ...)
- Start small, build out. Don't boil the ocean.
- Build, measure, iterate. Create tight develop-test loops with clear success criteria.



© 2025 Epic Systems Corporation. All rights reserved. Subject to [Terms of Use](#).

After Visit Summary, ASAP, Aura, Beacon, Beaker, Beans, BedTime, Best Care Choices for My Patient, Bones, Break-the-Glass, Buggy, Caboodle, Cadence, Canto, Care Everywhere, Charge Router, Cheers, Chronicles, Clarity, Cogito ergo sum, Cohort, Comfort, Community Connect, Compass Rose, Cosmos, Cosnome, Cupid, Discovery, Epic, EpicCare, EpicCare Link, Epicenter, EpicShare, EpicWeb, Epic Earth, Epic Nexus, Epic Research, Garden Plot, Grand Central, Haiku, Happy Together, Healthy Planet, Hello World, Hey Epic!, Hyperdrive, Hyperspace, Kaleidoscope, Kit, Limerick, Lucy, Lumens, MyChart, Nebula, OpTime, Phoenix, Powered by Epic, Prelude, Radar, Radiant, Resolute, Revenue Guardian, Rover, Share Everywhere, SmartForms, Sonnet, Stork, System Pulse, Tapestry, Trove, Welcome, Willow, Wisdom, With the Patient at Heart, and WorldWise are registered trademarks, trademarks, or service marks of Epic Systems Corporation in the United States of America and/or other countries. Other company, product, and service names referenced herein may be trademarks or service marks of their respective owners.

Patents Notice: [www.epic.com/patents](http://www.epic.com/patents).