

AI in Healthcare

Value vs. Hype

A Practical Framework for Evaluating AI in Your Practice

David Eitel, RRT · MHA · MSRT · RRT-ACCS
Senior Clinical Informatics Analyst — Intermountain Health



PRESENTED BY
HEALTHCARE UNCHARTED

1.0 AMA PRA Category 1 Credit™
Arizona Medical Association



Conflict of Interest Disclosure



No relevant financial relationships

No vendor payments, consulting arrangements, or speaker bureau relationships.
Employer: Intermountain Health (non-commercial health system)

Per ACCME Standards for Integrity and Independence in Accredited Continuing Education

About This Session

- 30-minute evidence-based session + Q&A
- All clinical claims cite peer-reviewed literature or primary source data
- No commercial support accepted or solicited



ARMA
ARIZONA MEDICAL
ASSOCIATION

Learning Objectives

1

Evaluate AI tools through workflow integration, data integrity, human oversight, and end-user experience

2

Apply a 5-question decision framework to assess AI vendor claims and distinguish evidence from hype

3

Recognize the role of physician leadership in AI governance and advocacy

4

Formulate concrete next steps to engage with AI evaluation at your institution or practice

Healthcare Is Not a Monolith

AI readiness depends on which part of healthcare you're talking about.



The Delivery of Healthcare

Operations, Administration, Workflow

- Scheduling, billing, prior authorization
- Clinical documentation & note-writing
- Supply chain, staffing, resource allocation
- Claims processing & revenue cycle

AI is already here now.

Bounded rules, high volume, low clinical risk, measurable outcomes.



The Practice of Medical Science

Diagnosis, Treatment, Clinical Judgment

- Diagnostic reasoning & differential diagnosis
- Treatment planning & patient-specific decisions
- Shared decision-making with patients
- Nuanced clinical judgment across contexts

AI is not here yet.

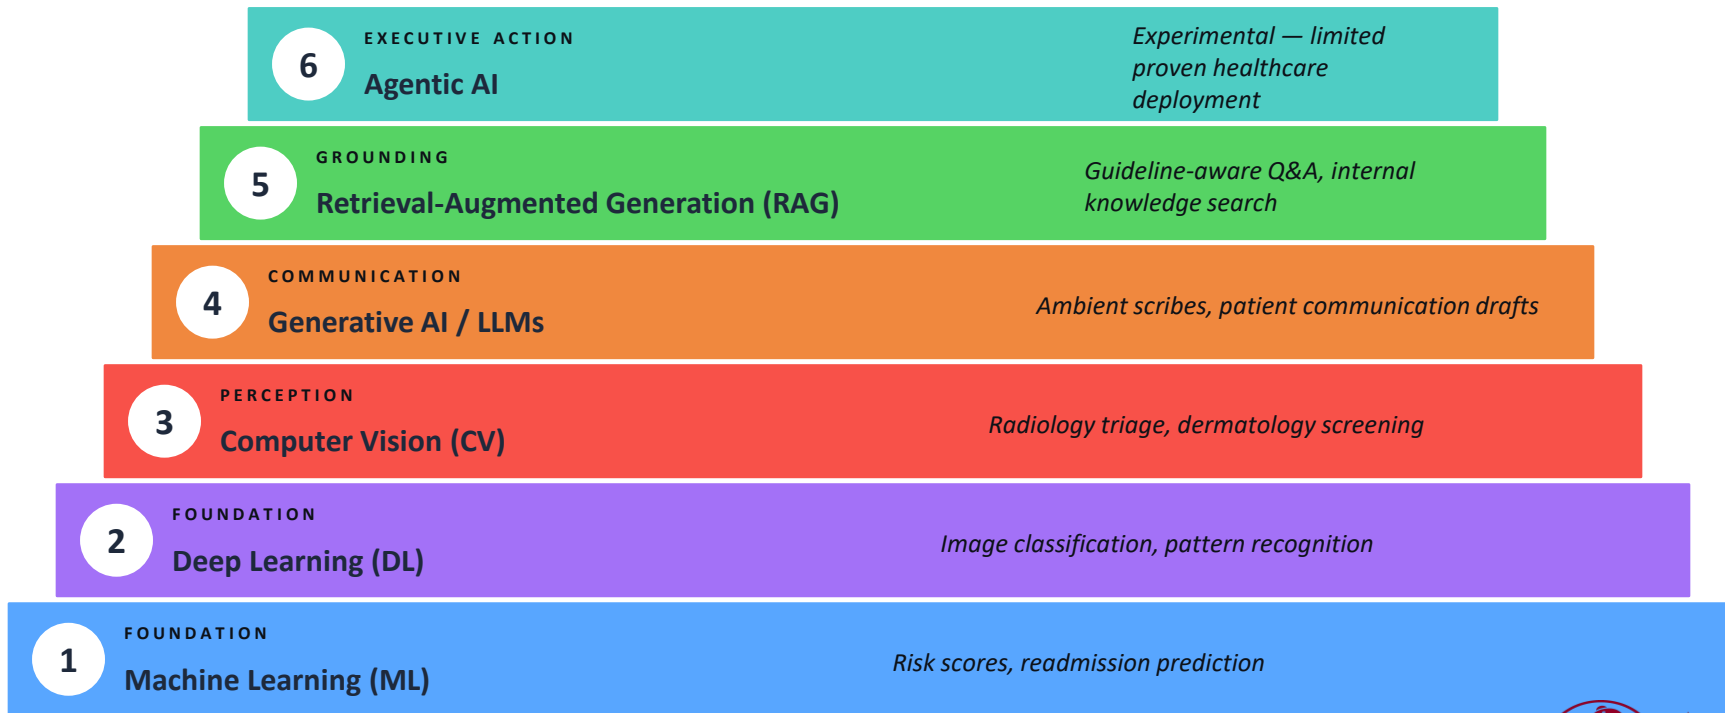
High variability, requires context, stakes are patient-to-patient.

Most AI failures happen when tools built for operations get deployed as if they're practicing medicine.



Not Everything Called "AI" Is the Same Thing

Six layers stacked on each other. Each does something different. Each carries different risk.



When someone says "we're using AI" — ask which layer. The answer changes everything.



AI in Healthcare by the Numbers

1,451

FDA-cleared AI devices (end 2025)

295

New clearances in 2025 alone

80%+

Physicians now using AI (AMA 2026)

The Growth Trajectory

2023

38%



2024

66%



2026

80%+

AI Is Already in Your Workflow

What's Running Now

- EHR-embedded risk scores & predictive models
- Sepsis prediction alerts (e.g., Epic Sepsis Model)
- Clinical Decision Support (CDS) tools
- Imaging triage and prioritization
- Ambient documentation scribes

FDA Shift: January 2026

CDS (Clinical Decision Support) software providing recommendations a clinician can independently review now falls outside FDA device oversight

More AI tools entering your workflow with less regulatory review

The burden of evaluation shifts to you and your institution



The question isn't whether to adopt AI — it's whether you know what you're already using.

ECRI's #1 Patient Safety Concern for 2026



"Navigating the AI Diagnostic Dilemma"

Named #1 patient safety concern by ECRI, the nation's leading patient safety organization



66%

of critical conditions missed by ML models in synthesized cases



Accuracy drops when prompts shift from textbook to real patient conversations

AI isn't foolproof. It can contribute to diagnostic mistakes for which you can be held liable.



The Epic Sepsis Model

0.63

AUC
(vs. 0.76 claimed)

33%

Sensitivity — missed
67% of sepsis cases

100s

Hospitals deployed
before external validation



Why This Matters

This is your EHR vendor's own AI, running in hospitals right now

Deployed nationally before independent external validation

2025 multicenter study (4 health systems) confirmed persistent performance issues in v2

Proprietary AI adopted at scale without external validation is a patient safety risk



The Bias Problem That Follows Every Tool

Optum/UnitedHealth Algorithm

Used on 200 million+ Americans annually

Used healthcare cost as a proxy for health need

Systematically deprioritized Black patients who were equally sick

Fixing the proxy: care for Black patients 17.7% → 46.5%

**The algorithm worked exactly as designed.
The bias was in how the problem was framed.**

Why This Repeats

Proxy bias: neutral variables encode structural inequality

Historic data inherits historic disparities

This will repeat with every tool unless you ask the right questions about training data

AMA AI Evaluation Guide, Domain 2: "Does the data reflect your patient population?"

The Liability Gap



Liability for Using Bad AI

If AI gives an incorrect recommendation and a physician follows it, the physician may be liable for not catching the error.



Liability for NOT Using AI

As AI becomes standard of care, physicians may face claims for NOT using available tools that could have caught what they missed.

Johns Hopkins researchers: asking physicians to judge when to trust AI is an "almost superhuman burden."

No major AI malpractice cases yet — but standard of care is shifting

Documentation + governance = your protection

Questions about your malpractice exposure? Contact MICA or your preferred carrier — before adopting any AI tool.

The Regulatory Patchwork



Federal (FDA)

1,451 cleared devices. Jan 2026: relaxed CDS oversight. New AI transparency/labeling guidance. No comprehensive national AI healthcare law.



States

250+ bills in 2025 across 47 states, 33 became law. Arizona HB 2175 — one of the first laws of its kind in the country — bans AI from denying claims without physician review (effective June 30, 2026). ArMA supported directly. Watch also: insurers are already using AI for downcoding; AZ HB 2407 addressed this but did not pass — it remains on the 2026 legislative agenda.



International

EU AI Act high-risk obligations: August 2026. CPT 2026 code set: 288 new codes for digital health & AI. Affects any device marketed globally.

No single rulebook. Your institution's governance IS the regulation for most AI in your workflow.

AMA Frameworks for AI Governance



8-Step Governance Toolkit

Released August 2025 | Manatt Health

For your institution:

Executive accountability & structure

Vendor evaluation process

Policy development & oversight

Available for CME credit



AI Tool Evaluation Guide

Updated Feb 2026 | 21 Specialty Societies

For evaluating any AI tool:

5 evaluation domains (next slide)

"AI Nutrition Labels" / model cards

Specialty-specific examples

FDA regulatory pathway guide

AMA: "Physicians must be full partners at every stage of the AI lifecycle."

Reading an AI "Nutrition Label"

01

Clinical Use Case & User

What does it do, who uses it, in what setting?

02

Training & Validation Data

Does the data reflect your patients and practice?

03

Risks & Mitigation

What can go wrong and how is it managed?

04

Effectiveness & Performance

How well does it work, validated externally?

05

Workflow Integration & Monitoring

Does it fit your workflow? Who monitors over time?

Cross-Cutting Considerations

Representation

Transparency

Clinical Relevance

Ongoing Safety

These 5 domains map directly to the 5-question framework you'll take home today.



The Pattern

Every AI tool that's working shares three traits:

1

Solves One Bounded Problem

Specific, well-defined task with measurable clinical outcomes.



2

Keeps a Human in the Loop

Physician reviews, approves, or overrides. Final decisions stay with a clinician.



3

Runs on Representative, Validated Data

Trained on data reflecting actual patients. Externally validated. Locally tested.



If a proposed AI tool doesn't meet all three: ask why not.

Evidence-Based Wins



Ambient AI Documentation

51.9% → 38.8%

Burnout reduction in 30 days (263 physicians, 6 systems)

Olson et al., JAMA Network Open, Oct 2025



Sepsis Early Warning (Bayesian)

10x fewer

False alerts. 46% more cases caught. 7x flagged pre-antibiotics

Cleveland Clinic / Bayesian Health, Sep 2025



Prior Authorization AI

60%

Prior auth processed in < 2 hours. Est. 5x ROI

Industry reports 2024–2025

All demonstrate the pattern: bounded problem, human in the loop, validated data.



The 5-Question Framework

Your quick-check before any AI adoption decision

1

What specific, bounded problem does this solve?

If the answer is vague, stop there.



2

Where does the physician stay in the loop?


The final decision must remain with a clinician.



3

What was this trained on, and does it reflect our patients?

Local validation is not optional.



4

How were clinicians involved in building this?

Not just consulted. Actually involved.



5

What does failure look like, and who's accountable?

If no one can answer this, that's your answer.



Apply this now: Think of one AI tool you've encountered. How does it score?



Your Action Steps

Everything here can be started without a budget, IT approval, or anyone's permission.

Ask for your institution's AI inventory

30 days

Download the AMA AI Tool Evaluation Guide

This week

Apply the 5-question framework to one tool

30–60 days

Advocate for an AI governance committee

60–90 days

Attend an AI demo — bring the 5 questions

Next available

Ask your EHR vendor what AI is already running

30 days

Take-home framework distributed at session close

References & Evidence Base

PEER-REVIEWED

- Olson A, et al. JAMA Network Open. Oct 2025. DOI: 10.1001/jamanetworkopen.2025.34976
- Wong A, et al. JAMA Internal Medicine. Jun 2021. DOI: 10.1001/jamainternmed.2021.2626
- Goh E, et al. JAMA Network Open. Oct 2024. DOI: 10.1001/jamanetworkopen.2024.40969
- Obermeyer Z, et al. Science. Oct 2019. DOI: 10.1126/science.aax2342
- Ostermayer D, et al. JAMIA Open. Dec 2024. DOI: 10.1093/jamiaopen/ooae133

ORGANIZATIONAL

- ECRI Top 10 Patient Safety Concerns 2026 (ecri.org)
- AMA Physician AI Sentiment Survey 2024, 2026 (ama-assn.org)
- AMA AI Specialty Collaborative: AI Tool Evaluation Guide, Feb 2026
- AMA STEPS Forward: Governance for Augmented Intelligence Toolkit, Aug 2025
- Cleveland Clinic / Bayesian Health Sepsis AI Pilot, Sep 2025

REGULATORY & POLICY

- FDA AI/ML-Enabled Medical Device List (Dec 2025 update)
- FDA Clinical Decision Support Guidance (Jan 6, 2026)
- Arizona HB 2175 — AI in Medical Claims (signed 2025, effective Jun 30, 2026)
- Manatt Health: Health AI Policy Tracker (manatt.com) — updated weekly



Getting Started with AI in Healthcare

Free resources for physicians at every level of AI familiarity

EVALUATION & GOVERNANCE

AMA AI Tool Evaluation Guide

5-domain framework, model cards

ama-assn.org

AMA STEPS Forward Governance

8-step framework, CME credit

ama-assn.org/steps-forward

AMA Center for Digital Health

Model policies, governance resources

ama-assn.org

EDUCATION & LITERACY

JAMA AI Special Report 2025

Summit consensus document

jamanetwork.com

ECRI Top 10 Safety 2026

#1: AI diagnostic dilemma

ecri.org

FDA AI/ML Device Database

All FDA-authorized AI devices

fda.gov

POLICY & TRACKING

Manatt Health AI Tracker

Weekly state & federal AI legislation

manatt.com

Arizona Medical Association

HB 2175, AI governance advocacy

azmed.org

Coalition for Health AI

Multi-stakeholder quality framework

coalitionforhealthai.org

FOR YOUR PRACTICE

Innolitics AI/ML Roundup

Annual FDA clearance analysis

innolitics.com

The Imaging Wire

Radiology AI tracking & analysis

theimagingwire.com

Healthcare Brew: AI

Accessible AI healthcare journalism

healthcare-brew.com

All resources are free. Photograph this slide or ask for the take-home PDF.





Questions & Discussion

What AI tools are you evaluating at your institution?

What governance gaps are you navigating?

Where does your skepticism need more evidence?

David Eitel, RRT · MHA · MSRT · RRT-ACCS
Senior Clinical Informatics Analyst — Intermountain Health | Adjunct Faculty – Utah Valley University

