

AI in Healthcare: The RT Reference Card

Operational Value, Limitations & Oversight · David Eitel RRT · MHA · MSRT · RRT-ACCS

THE CORE DISTINCTION

Operations AI vs. Clinical AI

OPERATIONS AI

Scheduling & staffing
Billing & coding
Prior auth processing
Ambient documentation

CLINICAL AI

Vent weaning decisions
Extubation prediction
Aerosol dosing adjustment
Deterioration alerts

Why it matters: Lower risk · already deployed vs. higher scrutiny · validation required. When someone says "AI in healthcare" — ask which column first.

THE FRAMEWORK

Evaluate Any AI Tool in 3 Questions

- 1 Which layer does this touch?**
Operations, Clinical, Financial, or Regulatory? The layer determines risk level and oversight requirements.
- 2 What happens when it's wrong?**
A scheduling error and a vent weaning error are not the same failure mode. Consequence shapes oversight needed.
- 3 Who owns the outcome?**
The vendor contract says it's not them. Your RT credential carries the responsibility — it doesn't transfer to the algorithm.

Use it: Put these three in your back pocket for every vendor demo, department pilot, and governance meeting.

WHERE IT'S WORKING

High-Confidence Operational AI

- Revenue Cycle & Prior Auth**
Coding accuracy, denial prediction. Real ROI, already scaled at major health systems.
- Ambient Documentation**
40-50% charting time reduction. DAX Copilot, Abridge, Suki. Workflow capture — not clinical judgment.
- Scheduling & Staffing**
Demand forecasting, patient flow modeling. Defined success metrics.
- Supply Chain & Equipment**
Inventory prediction, maintenance scheduling. High-volume structured data — AI's sweet spot.

Common thread: High-volume, structured, repeatable tasks with defined success criteria — very different from clinical judgment.

RED FLAGS

Warning Signs in Any AI Pitch

- Clinical outcomes promised with operational timelines ("90-day go-live")
- "Human in the loop" — with no workflow design to back it up at volume
- No clear answer when you ask: *Is this FDA SaMD cleared?*
- Validated on a different institution or population — not yours
- Vendor contract disclaims all clinical liability
- Alert volume designed to impress, not actionable at the bedside

The warning sign: Clinical outcomes promised with operational ease and a 90-day timeline — that gap is the risk.

INSURANCE AI: ALREADY HERE

What's Reviewing Your Prior Auths

300K

Claims denied by Cigna AI in 2 months (2022)

1.2s

Average algorithm review time per claim

90%

nH Predict denial rate in some Medicare cohorts

- ! Prior Auth Runs on AI**
Algorithms screen requests against coverage criteria before any human sees them. Denials are auto-generated.
- Document for the Algorithm**
Include explicit SpO₂ values, objective thresholds, criteria language. Under 2 seconds of review means keywords matter.

Why it matters: Knowing this changes how you write prior auths and appeal strategy. That is a clinical skill in 2026.

WHAT TO WATCH

Three Priorities for RT Practice

HIGH CONFIDENCE

- Ambient Documentation**
RT-specific tools are 18-24 months out. Ask your institution: are we piloting this?
 - Revenue Cycle AI**
Disrupting the prior auth layer now. Directly affects RT authorization workflows.
- VALIDATE FIRST**
- Predictive Deterioration**
Directly RT-relevant — demand population-specific validation data. Be in the governance room as an evaluator, not just the training session.

Your role: When your institution evaluates clinical AI, an RT should be in that evaluation room. Your skepticism is professional judgment.

HEALTHCARE UNCHARTED

Critical analysis of healthcare AI — for clinicians, by clinicians.

HealthcareUncharted.com

David Eitel RRT · MHA · MSRT · RRT-ACCS
MSRC Conference · April 2026 · Kalispell, MT