

Beyond the Buzz: Making AI Work in Real-World Medical Device Sales Training

Noah Zandan | CEO & Co-Founder, Quantified

TT LIFE SCIENCES • APRIL 15, 2026 • SALES TRAINING TRACK

20 MIN PRESENTATION • 10 MIN Q&A



Built With AI. Transparent About It.

Claude Cowork: Sonnet 4.6 Outline, draft copy, and speaker notes	OpenAI 5.4 Thinking Feedback on presentation. Fact-checking statistics and claims	Gamma Presentation Design
Instaheadshots Headshot	Flux Pro Images	Human Override Case Studies

📄 Transparency first: AI assisted every phase of this presentation. The judgment, framing, and decisions remained human.

If I'm going to ask you to trust AI in your field force, you should know exactly how I use it myself.



Most of you have already deployed something. Here's what's going wrong.

Honest State of Sales Training: 2026

Let's get pragmatic. What's real, what's hype, and what's coming.

✓ Working Now

- AI roleplay for objection handling & certification, deployed at scale in regulated environments
- Compliance scoring with auto-fail for non-compliant responses
- Rep-level skill-gap analytics surfacing specific coaching gaps
- Onboarding and certification cycle compression: 6-week programs running in under 2 weeks

🚫 Overpromised

- Fully autonomous sales reps replacing field activity
- Fully agentic AI solutions that are ready to deploy into sales workflows
- Plug-and-play deployment without change management
- AI that auto-adapts to your compliance framework out of the box

🕒 Watch: 12–18 Months

- Real enterprise AI agents that are valuable, observable, and safe (not just workflow automations)
- Stress-test messaging before launch, before a single rep is trained
- Adaptive difficulty adjusting in real-time based on personalized performance gaps

Three Failure Modes, And How to Detect Them Early

Not what goes wrong, but rather **how to catch it before it costs you.**

Compliance Theater

Signal: Tool scores confidence & keywords. A fluent off-label answer still passes.

Diagnostic: "What happens when a rep makes an off-label claim mid-simulation or can't answer a procedural question under clinical pressure?"

Intervention: Run a live scenario with a deliberate off-label answer using your content. If they pass, the evaluation is broken. In devices, a rep who can't answer a surgeon's clinical question mid-procedure doesn't get a second visit. The failure shows up in the account, not the training report.

The Adoption Cliff

Signal: Strong month-one completion. By month three, reps do the minimum and no more.

Diagnostic: "What sessions are voluntary vs. mandated, and how does that trend over 90 days?"

Intervention: Measure voluntary return rate in week two of your pilot. Low return = realism problem.

The Dashboard Mirage

Signal: Completion, logins, time-on-platform, but you can't tell your VP which reps are ready.

Diagnostic: "Can you tell me this rep is ready to call a cardiologist about this indication next Tuesday?"

Intervention: Define "certified" in performance terms before deployment. Completion and readiness are different data.



The Human-AI Balance: What Leaders Get Wrong

"This exists so your manager's time goes to the conversations that matter, not watching you certify something a computer can assess."

AI Handles

- Consistent on-demand repetition at scale
- Objective scoring against fixed standards
- Surfacing skill gaps across large cohorts
- Certification without trainer bottlenecks
- Longitudinal performance tracking

Humans Must Own

- Reading emotional state and adjusting
- Contextual coaching from field reality
- Rep relationships that drive accountability
- Readiness calls data alone can't make
- Motivating a rep struggling with confidence

❏ **The mistake:** Framing AI as a replacement for human coaching. Managers feel displaced. Adoption fails before the first session runs. Not because the tool is bad, but because the framing created resistance.

The Data Security Question Nobody Asks Early Enough

The Core Question to Ask Every Vendor

"Where does content live at the end of a session, and can it influence your model?"

Most platforms train on user inputs by default. Pre-launch clinical data, competitive positioning, and investigational device messaging become potential training material. Discovery typically comes from legal review *after* deployment.

Get These Four Commitments in Writing

01

Model Training Clause

Customer content is NOT used to train, fine-tune, or update any model, in the agreement.

02

Data Isolation

Content processed in an isolated environment and does not persist after session completion.

03

SOC 2 Type II

Certification is current and available for review on request.

04

Audit Logs

Session-level audit logs available on demand, no reconstruction required.

What Compliant Training Actually Requires

Three non-negotiables, and the diagnostic question to test each one.

1

Score Against *Your* Messaging

Your approved messaging, your indications, your specific language, not a model rubric that approximates your standards.

Ask: Does it score against your PI, your detail aid, your approved objection responses?

2

Non-Compliant Answers Auto-Fail

Automatic, consistent, every session. A rep who makes an off-label claim passes or fails based on the system, not who reviewed the session.

Ask: Show me what happens with a deliberately off-label answer, live, using your content.

3

Audit Trail Built from Day One

Every session logged, every score defensible, no reconstruction required. The trail exists before you need it, not because something went wrong.

Ask: Can you produce session-level logs on demand, without requesting a data export?

Five Things I'd Do Starting From Scratch Today

In order of importance, the decisions that determine whether this succeeds or stalls.

1 Audit your current certification process first

Count actual hours: manager time, rep time, repeat attempts, admin overhead. That number becomes your business case and your baseline.

2 Run one pilot: one product, one cohort

Define success before you start. Not completion rate – readiness. What does a rep who passes do differently in front of a customer?

3 Require the three compliance non-negotiables in writing

Scoring against your approved messaging, auto-fail for non-compliant answers, audit-ready logging. In the contract, not the sales deck.

4 Ask where your content lives before legal does

"Does my content influence your model?" Simple question. The answer tells you everything, get it in writing before your internal review.

5 Measure voluntary return rate in week two of your pilot

The earliest leading indicator of simulation realism. Mandated completion tells you nothing. Unprompted return tells you everything.

Three Questions to Take Back With You

"The tool is the easy part. The requirements are the work."

1

Off-Label Test

If a rep answers off-label in your training tool today, what actually happens?

2

Certified Means What?

What does "certified" mean for your team, in performance terms, not completion terms?

3

Actionable Data?

Does your training data give front-line managers something to *act on*, not just something to report?



The takeaway: The teams getting value from this aren't asking whether AI works for training. They're defining what ready means and measuring backward from there.

Further Reading + Connect

Recommended Reading

One Useful Thing – Ethan Mollick

Newsletter: The best ongoing source for practical, grounded AI thinking

Superagency – Reid Hoffman

What could go right when humans partner with AI at scale

Co-Intelligence – Ethan Mollick

Living and working with AI – a framework for human-AI partnership

The AI-Driven Leader - Geoff Woods

Harnessing AI to Make Faster, Smarter Decisions

Find Me

Noah Zandan CEO & Co-Founder, Quantified

noah@quantified.ai

quantified.ai



Stop by the **Quantified booth**, happy to keep the conversation going.

 Quantified