

Healthcare-specific AI

What it is, why it matters and how to know if you have it

Meet your hosts.



Dale Van Gorder
Chief Sales Officer

Brings the market perspective — what organizations are experiencing with AI in the field today.



Lily He
Chief Product Officer

Brings the product and technical depth — how the system is built and why the design choices matter.

The AI moment is already here.

Healthcare organizations are investing in AI at a pace that outstrips their ability to govern it.

\$2T+

Global AI spending forecast by
2026

*Gartner, 2025

78%

of organizations report using AI
in at least one business
function

*McKinsey State of AI, 2025

37%

YOY growth in AI investment,
the fastest in history

*Gartner, 2025

The question isn't whether AI is coming to healthcare, it's whether healthcare is ready for it.

AI is already in healthcare, but governance is not keeping up.

#2

Patient safety threat in 2025

"Insufficient governance of artificial intelligence" in healthcare.

*ECRI Top 10 Patient Safety Concerns, 2025

16%

of hospital executives have a systemwide governance policy for AI use and data access.

*2023 hospital executive survey

63%

of organizations don't have, or aren't sure they have, the right data management practices for AI.

*Gartner, Q3 2024

Not all AI is the same. And not all of it knows healthcare.

HIGH RISK

General-Purpose LLM

Trained on broad internet data. Answers quickly. Knows a little about everything. Knows nothing about your organization.

MODERATE RISK

Retrieval-Augmented Generation (RAG)

Pulls from a knowledge base before answering. Better than pure LLM but only as good as the documents it retrieves.

VARIABLE RISK

Fine-Tuned Domain Model

Trained on domain-specific data. More relevant outputs, but studies show fine-tuning can actually increase hallucination rates.

LOWEST RISK

Healthcare-Specific + Human Oversight

Built on healthcare data. Validated by human experts. Outputs checked before use. Errors caught before they compound.

"Strikingly, general-purpose models achieved higher proportions of hallucination-free responses than medically-specialised models — suggesting narrow fine-tuning may actually increase certain hallucination risks."

*MedRxiv / Frontiers in AI, 2025

The problem isn't just that AI gets it wrong. It's that AI gets it differently every time.

Low accuracy · Low precision



High accuracy · High precision



 **NTRACTS**

Repeatability of results isn't a nice-to-have. It's the standard.

Ntracts contract data abstraction is both accurate and precise.

Two types of AI output

Deterministic


There is a right answer, and the system must find it consistently. Ntracts contract data abstraction works this way. An expiration date is either correct, or it isn't.

Non-deterministic

Outputs are variable by design. General purpose LLMs work this way. Ask the same question three times, get three different answers.

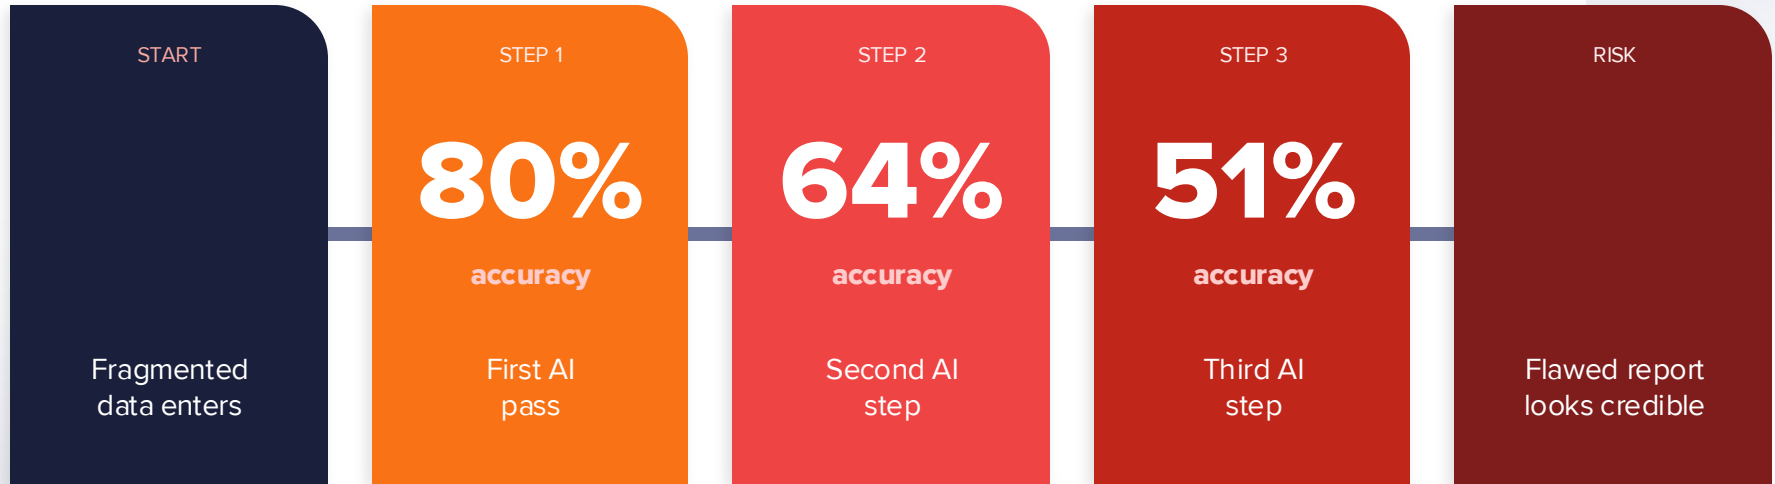
AI doesn't know when it doesn't know.

Hallucination rates across domains, when AI generates confident, plausible, wrong answers.



80% accuracy isn't 80% risk. It's 100% risk.

Each AI step that uses flawed data multiplies the error, and the output still looks credible.



$80\% \times 80\% = 64\%$ | $64\% \times 80\% = 51\%$ | By the time the third AI step completes, you may be operating at barely half accuracy.

And by that point, the output looks completely credible. The error is invisible.

Governments aren't waiting.



Neither should you.

State legislatures are already passing laws requiring human oversight of AI in healthcare decisions. The patchwork is growing fast.

TechTarget / State Health AI Laws, 2026

60+

Bills targeting payer AI introduced in 2025

4 enacted into law in Arizona, Maryland, Nebraska and Texas prohibiting AI alone from denying care and requiring human review of algorithm-driven decisions.

TraxTech / State AI Regulation Report, 2025

5

States limiting AI-driven coverage denials

Arizona, Connecticut, Maryland, Nebraska and Texas have passed legislation prohibiting insurers from using AI as the sole basis to deny medical care coverage.

Manatt Health AI Policy Tracker, 2026

7

States targeting AI downcoding in 2026

California, Connecticut, Illinois, Indiana, Maryland, Missouri and Oregon introduced bills prohibiting AI from downcoding claims without oversight by a licensed physician.

California SB 1120, enacted 2024

CA

SB 1120 — Physicians Make Decisions Act

Restricts health plan utilization-review AI from supplanting provider judgment. Mandates human processes and documentation for all AI-driven coverage decisions.

If governments are requiring human oversight of AI decisions, your organization's AI governance standard should be no different.

Poor data quality is already costing organizations. AI makes it worse.

95%

of organizations say poor data quality negatively impacts business outcomes

*Experian, 2022 Global Data Management Report

\$12.9M

lost per year on average due to bad data quality

*Gartner Research

20%

decrease in productivity caused by poor-quality data, plus a 30% increase in costs

*McKinsey Global Institute

AI spending is forecast to surpass \$2 trillion in 2026. When AI investment scales, the cost of poor data quality scales with it.

Most AI initiatives never make it past the pilot.

The reason is almost always data, not technology.

60%

of AI projects will be abandoned through 2026 due to poor data quality

*Gartner, 2024

95%

of generative AI pilots fail to progress beyond experimentation

*MIT NANDA Study

30%

of GenAI projects abandoned after proof of concept by end of 2025

*Gartner, 2024

16%

of AI initiatives have successfully scaled across the enterprise

*IBM Institute for Business Value, 2025

Knowing clean data matters is step one. Making it happen is step two.

 **NTRACTS**

Not all data supports compliant AI.

The difference is how your data is structured, validated and maintained.

UNSTRUCTURED DATA

Common but difficult to rely on.

- ✗ Policies, contracts and audits stored as narrative text
- ✗ Key terms buried and inconsistently interpreted
- ✗ Limited traceability across workflows
- ✗ Requires multiple AI steps, each one adding error
- ✗ No audit trail from source to output

STRUCTURED DATA

Purpose-built for compliance and AI.

- ✓ Key terms captured as discrete, trackable elements
- ✓ Aligned to compliance workflows and reporting
- ✓ Clear audit trail from source to outcome
- ✓ Fewer AI steps, improving overall accuracy
- ✓ Defensible in audits and regulatory reviews

Healthcare can't afford to be almost right.

In general business, "almost right" means:

- A report that needs a revision
- A recommendation worth questioning
- A process that runs slower than expected
- A decision that can be walked back

In healthcare, "almost right" means:

- A compliance failure an auditor finds first
- A policy that guides clinical practice incorrectly
- A contract that auto-renews on terms no one reviewed
- A billing pattern that triggers a Medicare investigation

Would you hand an AI-generated policy draft to a surveyor?
Would you defend it after an incident? If you hesitated — that's the gap.

AI should reduce your risk and free your team to do the work only they can do.



Most organizations think about AI as a time-saving tool.
In healthcare, the real outcomes are more specific than that.

Reduce organizational risk

AI that is healthcare-specific, built on clean data and validated by humans surfaces problems before they compound.

- Compliance gaps caught early, before they reach an auditor.
- Contract exposure visible before renewal, not after.
- Policy gaps flagged before a survey, not during one.

Work to the top of your license

Healthcare professionals were hired for judgment, not administrative logistics. The right AI removes the work that doesn't require their expertise.

- Compliance teams focused on risk strategy, not inbox management.
- Contract teams reviewing terms, not chasing signatures.
- Clinical staff using policies, not maintaining them.

These aren't aspirational outcomes.
They're the standard healthcare organizations should be holding their AI to.

Five questions every healthcare organization should ask about their AI.

1 What data was this AI tested on?

Generic internet data? Industry-specific? Your organization's actual contracts and policies?

2 What happens when it's wrong?

Is there a human review step before output is used? Or does the AI output go directly into a workflow?

3 How is your data structured before it enters the AI?

Are your contracts and policies in narrative text, or structured, validated, tracked elements?

4 Can you explain an output to a surveyor?

Could you trace the AI's answer back to a source document and defend the logic?

5 Is accuracy maintained over time, or just at go-live?

Data changes. Contracts renew. Policies update. Does the AI's foundation stay current?

Clean data doesn't happen automatically. It's built.

A structured approach to clean contract data is the foundation Ntracts builds from day one of implementation.

01

Contract ingestion
& normalization

Raw documents brought in and standardized into consistent formats and structures.

02

Clause identification
& validation

Key terms identified, extracted and validated against healthcare-specific logic.

03

Metadata accuracy
& alignment

Dates, parties, obligations and terms captured as discrete, trackable data points.

04

Ongoing governance
& updates

Data maintained and updated as contracts change, accuracy preserved over time.

"Successful AI adoption starts with getting the data right during implementation and continues with keeping it accurate over time."

What makes Ntracts AI different.

Built for healthcare from the start. Validated by human expertise at every step. Grounded in structured, clean data.

Built for healthcare

Built specifically for healthcare use cases including details like PII abstraction that generic AI won't prioritize. We build toward the things that matter in healthcare, not the things that matter everywhere else.

Human oversight built in

AI supports your team — it does not replace it. Outputs are designed with expert input. Final validation stays with the people responsible for compliance. Nothing moves forward unchecked.

Structured, validated data

Data is structured from day one of implementation. Key terms captured as discrete, trackable elements. Consistent data foundation across contracts and reporting. Clean in — trustworthy out.

Nothing moves forward unchecked.

From AI analysis to human review — every output is validated before it becomes a decision.

AI Analysis

- Clause identification
- Risk flagging
- Data abstraction
- Summarization of results

+

Human Expertise

- Healthcare context
- Regulatory knowledge
- Validation & review
- Accountability

+

Accurate Outcomes

- Trusted reporting
- Defensible decisions
- Confident compliance
- Reduced risk

AI provides speed and scale. Human expertise provides judgment and accountability.
Together they create outcomes neither can achieve alone.

What Ntracts AI looks like in practice.

AI Reporting Assistant

Ask natural-language questions and receive dynamic reports instantly. No query-building required. The underlying query is always visible so your team can validate and adjust results.

Example: "Show me all contracts expiring in 90 days by department."

AI Clause Search

Search your entire contract library by meaning, not keywords. Surface relevant clauses in seconds and open full clause context to review language in place and act with confidence.

AI Contract Review

Review and negotiate contracts faster with AI-assisted redlining. Ntracts surfaces risks, suggests edits and supports more consistent contract language.

AI Contract Abstraction

Upload a contract and Ntracts automatically extracts key data elements: parties, effective and expiration dates, auto-renewal provisions, termination clauses and PHI access. Every suggested value is highlighted with source text, your team accepts or rejects. Nothing moves forward without human review.

AI Contract Abstraction

Upload a contract and Ntracts automatically extracts key data elements: parties, effective and expiration dates, auto-renewal provisions, termination clauses and PHI access. Every suggested value is highlighted with source text, your team accepts or rejects. Nothing moves forward without human review.

Policy Data Abstraction

The same structured data abstraction approach applied to contracts, applied to policy documents. Key policy elements extracted, validated and governed from day one so your policy library is searchable, current and audit-ready.

Where AI efficiency meets clinical expertise.

AI produces policy language quickly. RN expertise makes it clinically defensible.

Scenario	AI Draft	Enhanced by Ntracts RN	Real-World Impact
Behavioral health de-escalation	"Staff should respond to agitated patients using de-escalation techniques."	"Staff shall perform an immediate safety assessment, initiate trauma-informed de-escalation per the facility's behavioral response protocol, document all interventions and escalate to clinical leadership if risk increases."	AI offers general intent. RN-informed policy provides actionable, defensible steps.
Ambulatory medication reconciliation	"Medication reconciliation must be completed at each patient visit."	"The RN shall complete medication reconciliation upon admission, documenting medications by name, dose, route and frequency. The managing provider will review, resolve discrepancies and ensure changes are entered by the attending physician."	AI states the requirement. RN guidance defines ownership, timing and documentation.
Infection prevention in long-term care	"Facilities must follow infection control standards."	"Nursing staff shall implement hand hygiene audits, initiate transmission-based precautions when indicated, complete environmental cleaning checklists and review infection data weekly with leadership."	AI references standards. RN policy operationalizes prevention in real care environments.

Knowing clean data matters is step one. Making it happen is step two.

BAD DATA IN

Fragmented contracts
Inconsistent fields
Narrative text without structure

POOR INSIGHTS OUT

- ✗ Unclear reporting
- ✗ Questionable answers
- ✗ Unnecessary risk
- ✗ Compliance failures that look credible

VS

CLEAN DATA IN

Standardized contracts
Validated clauses
Structured, tracked data elements

TRUSTWORTHY INSIGHTS OUT

- ✓ Confident reporting
- ✓ Reliable answers
- ✓ Informed decisions
- ✓ Compliance outcomes you can defend

When the data foundation is right, everything that follows is more reliable.

Clean data and human expertise together are what make AI outputs in healthcare trustworthy, defensible and worth acting on.



Trusted reporting

AI-driven reporting becomes a source of truth rather than a point of skepticism. Leadership can rely on it.



Reliable clause search

Search across your entire contract library and get answers you can act on — not answers you have to verify manually.



Confident compliance

Audit readiness is continuous, not assembled the week before a site visit. Evidence is always current and traceable.



Faster decisions

When teams trust the output, decisions happen faster. Less time verifying, more time acting on what the data says.



Defensible AI

Every AI output can be traced to a source. You can explain it to a surveyor, a board or legal counsel — and stand behind it.



Reduced organizational risk

Compounding errors stop compounding. Flawed policies stop reaching clinical practice. Compliance gaps surface before audits do.

Thank you.

Three things to take away from today.

- 1 AI does not fix bad data. It amplifies it. Start with the foundation.
- 2 Errors compound with every AI step. 80% accuracy becomes 51% in three steps and the output still looks credible.
- 3 Responsible AI in healthcare requires structured data, healthcare-specific training and human oversight at every step.

See how Ntracts builds healthcare-specific AI.

Request a demo at ntracts.com | sales@ntracts.com | (888) 316-9805

Q&A

Questions welcome.