AI Compliance – How to Protect Yourself

Real-World Case Studies & Compliance Challenges

AR HFMA Spring Conference | 2026

Presented by AccuCodeAI

1

Over Three Short Years

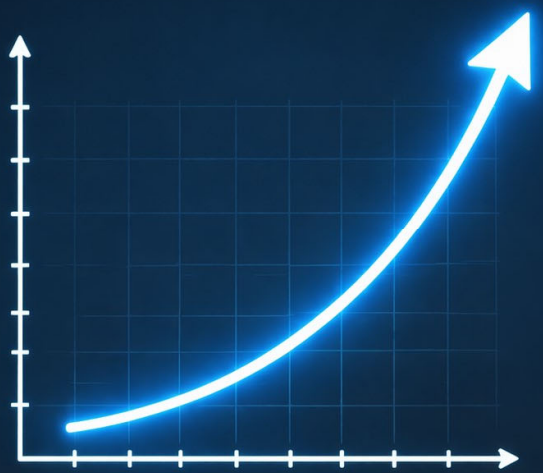

2

What We'll Cover Today

3

Artificial Intelligence Policy Landscape |

U.S. Congress · Executive Orders · Arkansas

U.S. Congress

119th Congress · 2025–2026

No AI Healthcare Bills have passed committee
Strategy to defend against financial crimes using AI, including fraud & misinformation (H.R. 2152)

H.R. 238 – The Healthy Technology Act of 2025
Billed proposed that AI or machine learning technology may qualify as a prescribing practitioner

H.R. 5045 – The HEALTH AI Act
A bill to direct the Secretary of HHS to establish a grant program to facilitate research regarding the use of generative AI in healthcare

NDAA FY2026
The only AI related bill to go into law. Requires Department of Defense AI governance policy, cybersecurity standards, and cross-functional oversight framework

Executive Orders

Trump Administration · 2025–2026

Jan 2025 — EO 14179
Removed Biden-era AI guardrails; redirected policy toward deregulation & U.S. global dominance

July 2025 – AI Action Plan
Released comprehensive national AI strategy; signed EO preventing 'woke AI' in federal government

Dec 2025 – National Framework EO
Created DOJ AI Litigation Task Force to challenge state AI laws; threatened to withhold BEAD broadband funding

March 2026 — Legislative Blueprint
White House released 7-area framework urging Congress to preempt state laws with a light-touch federal standard

Arkansas

State Action · AI CoE & the 95th General Assembly

Governor Sanders as an AI Leader
Challenged & Defeated components of the Big Beautiful Bill that would have limited State level AI protections

Executive Actions in the Governor's Office
Created an Arkansas AI taskforce through the AI & Analytics Center of Excellence – June 2024

The 95th General Assembly Failed to Pass AI Healthcare Legislation
All proposed bills were withdrawn before making it to a committee vote

The Arkansas Takeaway
Arkansas has made attempts to keep the ability to regulate Artificial Intelligence at the State Level but has not been able to pass legislation in a far less divided house and senate

1,000+ state AI bills introduced in 2025 · 38 states passed AI legislation in 2025 · Only 14% of health orgs have implemented AI tools · Arkansas recruiting first-ever Chief AI Officer (2026)

4

AI USE CASES ACROSS THE REVENUE CYCLE

From Patient Access to Analytics — Where AI Creates Value

FRONT END	MID CYCLE	BACK END	ANALYTICS
Eligibility Verification AI checks insurance coverage in real time at scheduling, reducing front-door denials.	**Autonomous Coding** NLP translates clinical documentation into ICD-10/CPT codes with stated 95% accuracy, no human touch, if done properly.	**Denial Prediction** Predictive models flag high-risk claims before submission, reducing denial rates by 20–30%.	**Fraud Detection** AI analyzes billing patterns to surface anomalies that may indicate fraud, waste, or abuse.
Prior Authorization NLP reads clinical notes and auto-submits auth requests, cutting days off approval timelines.	**CDI & Documentation** AI flags incomplete or vague documentation in real time, prompting clinicians before discharge.	**Denial Management** AI agents prioritize, route, and draft appeal letters for denied claims, accelerating resolution.	**Contract Modeling** ML simulates payer contract scenarios, identifying underpayments and optimizing renegotiations.
Patient Estimates ML predicts patient cost responsibility at point of scheduling to improve collections.	**Level-of-Care Scoring** Continuous AI scoring helps assign inpatient vs. observation status aligned to MCG/InterQual criteria.	**AR Prioritization** ML ranks accounts receivable by collectability, directing staff time to highest-yield accounts.	**Revenue Forecasting** Predictive analytics models future cash flow and denial trends based on historical RCM data.
Appointment Optimization AI matches scheduling to payer mix and service capacity to maximize reimbursable volume.	**Charge Capture** ML identifies missed charges by comparing clinical activity to billed items, closing revenue leakage.	**Payment Posting** RPA auto-posts EOBs and ERA files, reconciling payments and adjustments without manual entry.	**Voice AI / Ambient** Ambient voice capture auto-generates clinical documentation during patient encounters for billing.

AI use cases span every stage of the revenue cycle. Compliance governance and human oversight requirements apply across all phases.

5

The Compliance Landscape

Why AI + Revenue Cycle Creates Unique Regulatory Exposure

23%	$16B	46%	28%
Increase in claim denial rates 2016–2022	Annual revenue leakage from billing errors	Hospitals now using AI in RCM operations	Provider AI confidence (down from 68% in 2022)

False Claims Act	HIPAA & PHI	RAC/OIG Audits	Algorithm Liability
AI-generated billing errors can trigger federal FCA liability — no intent required.	NLP/LLM models that train on or process patient data require strict BAA governance.	Automated codes are subject to the same Recovery Audit scrutiny as manual codes.	Liability falls on the hospital — not the AI vendor — when a model produces errors.

6

AI & Revenue Cycle

COMPLIANCE CHALLENGES

Hospital Revenue Cycle Management in the Age of AI

72%

of executives prioritize AI/automation investment

HIPAA & Data Privacy

PHI access at scale creates breach exposure

Algorithmic Transparency

Regulators require explainable AI decisions

Regulatory Alignment

Models go stale when payer rules shift

AI Coding Errors & Upcoding

Systematic miscodes trigger False Claims risk

Human Oversight

90% of leaders say oversight is essential

Vendor Accountability

Third-party AI gaps create compliance holes

Data Quality & Integrity

Bad input data amplified across thousands of claims

Workforce Readiness

63% of RCM depts report critical staffing gaps

Sources: HFMA · Experian Health · Medical Economics · AAPC · EY Healthcare 2025

7

CASE STUDY 01 | HIPAA and Proper Use and Access Standards

Northwell Health | New York State's Largest Health System

22	800+	16K
Hospitals	Outpatient Sites	Physicians

Why Scale Amplifies Compliance Risk

With millions of documents to process, Northwell could not cover clinical NLP review manually. A 10-year agreement with Clinithink deployed CLiX NLP to automate chart reviews and denials management — but at this scale, a single systematic error in the model affects thousands of claims.

HIPAA & Data Governance Challenges

Business Associate Agreement (BAA)

Any NLP vendor processing PHI must be a formal HIPAA Business Associate. Failure to execute a BAA before data sharing begins is an immediate HIPAA violation.

Minimum Necessary Standard

NLP models must only access patient data necessary for the specific RCM function. Training models on broader clinical datasets without restriction violates the minimum necessary rule.

Access Controls & Audit Logging

All AI interactions with PHI must be logged, auditable, and restricted to authorized users — requirements that many legacy healthcare IT architectures were not built for.

Breach Notification Risk

If a vendor's AI infrastructure is breached and PHI was involved in model training or inference, Northwell — as covered entity — bears the breach notification obligation.

8

9

LLM TRAINING DATA & HIPAA COMPLIANCE

How Hospital Data Trains an AI Model — and Where the Risk Lives

HOW HOSPITAL DATA ENTERS AN LLM

1 Raw Data Extraction — Contains PHI

Clinical notes, EHR records, billing claims, coding histories, and denial letters are exported from hospital systems — often containing patient names, diagnoses, dates of service, and procedure details.

2 Data Preprocessing & Labeling — PHI Exposure Risk

Data scientists clean, tokenize, and label the records. Human annotators may review real patient encounters to create training examples. PHI is often still present at this stage.

3 Upload to Training Environment — Potential HIPAA Disclosure

Processed data is transferred to the vendor's cloud or on-premise GPU infrastructure. If a third-party vendor hosts the training environment, this transfer may constitute a HIPAA disclosure.

4 Model Fine-Tuning — PHI Embedded in Model

The LLM learns patterns from the training data. Because neural networks memorize training examples, patient-specific details can become embedded in model weights — recoverable via adversarial prompts.

5 Model Deployment & Inference — Ongoing Exposure Risk

The trained model is deployed for RCM tasks. Each inference call may expose training data via model memorization, even if no raw data is stored in the production environment.

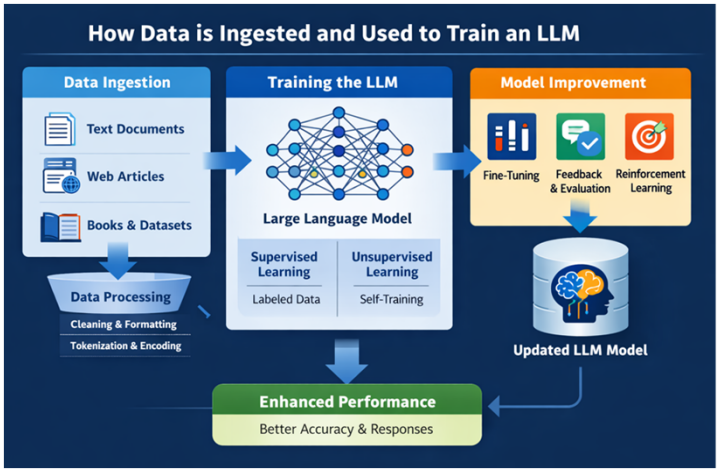

Key Regulation: 45 CFR §§ 164.502, 164.308, 164.312 (HIPAA Privacy & Security Rules) | OCR Guidance on Cloud Computing (2016) | npj Digital Medicine (2025)

10

LLM TRAINING DATA |

HIPAA Compliance Risks When Hospitals Use LLMs in Revenue Cycle

6 HIPAA RISK CATEGORIES — COVERED ENTITY LIABILITY APPLIES TO EVERY ONE

Missing or Invalid Business Associate Agreement (BAA) — CRITICAL

HIPAA requires a signed BAA before any vendor receives, processes, or stores PHI. Many AI vendors offering LLM fine-tuning will not sign a BAA — making any data sharing with them an automatic HIPAA violation regardless of outcome.

45 CFR § 164.308(b)

PHI Memorization Inside Model Weights — CRITICAL

LLMs can memorize and reproduce specific training records — including patient names, diagnoses, and billing data — when prompted. PHI does not disappear after training; it persists indefinitely in model weights and can be extracted via adversarial prompting.

45 CFR § 164.502

Data Transmitted to Third-Party Cloud Infrastructure — HIGH

Fine-tuning on a vendor's cloud platform (OpenAI, Google, Azure) transmits PHI outside the hospital's firewall. Without a HIPAA-compliant infrastructure agreement in place, this constitutes an unauthorized disclosure — even if data is later deleted.

45 CFR § 164.312(a)(2)

Human Annotators Reviewing Patient Records — HIGH

LLM data labeling often involves human reviewers — sometimes offshore contractors — reading real patient records. These individuals may not be covered by the hospital's workforce training or access controls, creating minimum-necessary and workforce security violations.

45 CFR § 164.308(a)(3)

Hospital Bears Full Liability as Covered Entity — CRITICAL

HIPAA places enforcement responsibility on the covered entity — the hospital — not the AI vendor. A breach caused by a vendor's LLM pipeline exposes the hospital to OCR investigation and civil monetary penalties up to $1.9M per violation category per calendar year.

45 CFR § 164.530(f)

No Right of Erasure: PHI Cannot Be Deleted from a Trained Model — MEDIUM

Unlike a database, you cannot delete a specific patient's record from trained neural network weights without retraining the model entirely. If a patient exercises HIPAA rights to restrict data use, the hospital may have no technical mechanism to comply once training is complete.

45 CFR § 164.522

Remember: HIPAA liability falls entirely on the hospital as the covered entity — not the AI vendor. Vendor contracts do not transfer this obligation.

Key Regulations: 45 CFR §§ 164.502, 164.308, 164.312, 164.522, 164.530 | HHS OCR Guidance on Cloud Computing (2016) | npj Digital Medicine (2025)

11

CASE STUDY 02

Inova Health System | Falls Church, Virginia

The Challenge

Rolling out multiple AI solutions at scale across a large provider organization while juggling product oversite and ongoing workflow challenges related to staffing and regulatory changes.

Inova struggled with manual workflows for utilization review and lacked visibility into observation rates and peer-to-peer utilization. Critically, misclassifying patients as inpatient vs. observation status exposes hospitals to CMS Recovery Audit Contractor (RAC) audits and Medicare repayment demands — a major compliance risk.

AI Solutions Deployed

- Nym autonomous coding engine for ED coding — $500K annual cost reduction
- Xsolis Dragonfly platform with AI-driven Care Level Score™ updating in real time
- Claims 80% of cases reviewed within first 24 hours of admission

Compliance Focus

RAC Audit Exposure

Incorrect inpatient vs. observation classification is one of the most audited billing issues under Medicare. AI must align with InterQual/MCG clinical criteria.

Documentation Integrity

AI Care Level Scores are only defensible if backed by complete clinical documentation. Gaps still trigger payer denials regardless of AI output.

Governance Gap

Mid-cycle and back-end denial teams historically worked at cross-purposes. Inova had to build new governance structures to align AI outputs across functions.

12

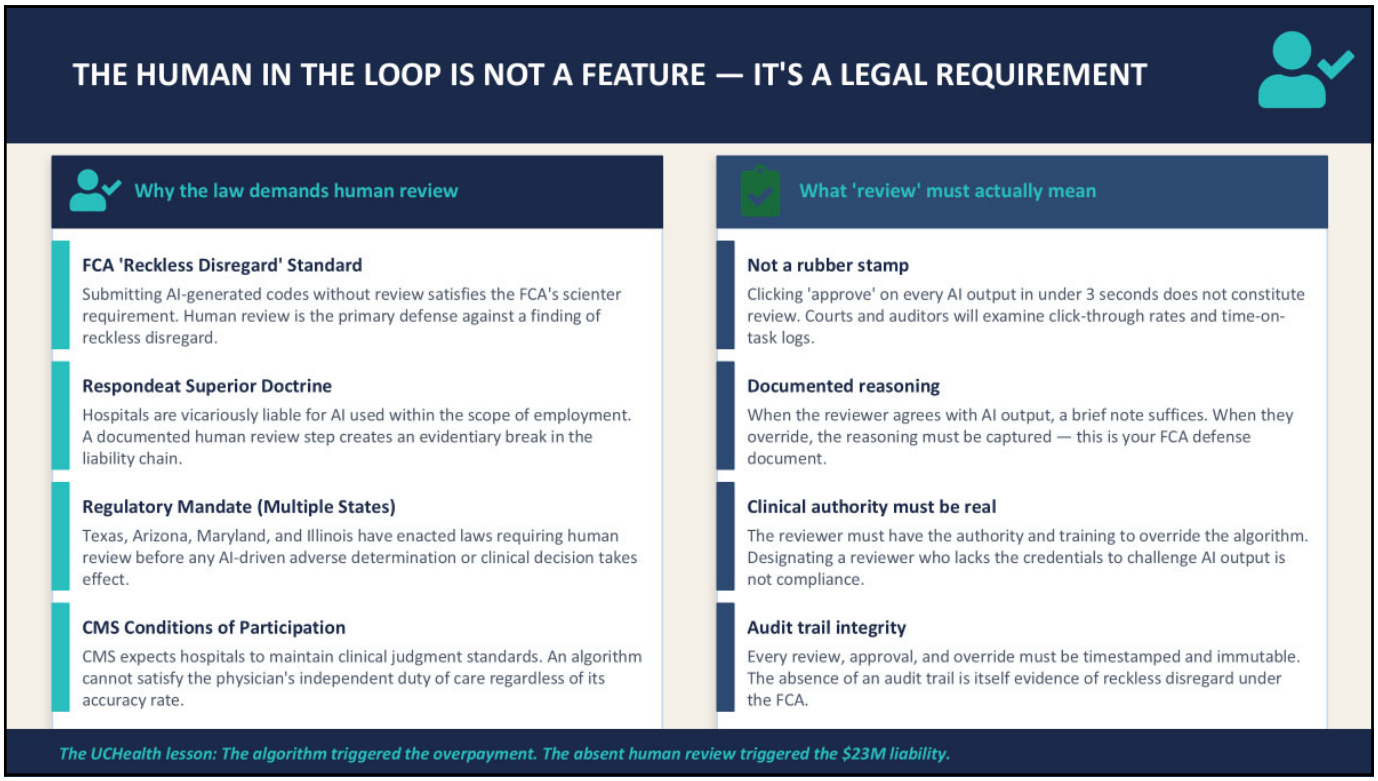
THE HUMAN IN THE LOOP IS NOT A FEATURE — IT'S A LEGAL REQUIREMENT
Why the law demands human review
FCA 'Reckless Disregard' Standard
Submitting AI-generated codes without review satisfies the FCA's scienter requirement. Human review is the primary defense against a finding of reckless disregard.
Respondeat Superior Doctrine
Hospitals are vicariously liable for AI used within the scope of employment. A documented human review step creates an evidentiary break in the liability chain.
Regulatory Mandate (Multiple States)
Texas, Arizona, Maryland, and Illinois have enacted laws requiring human review before any AI-driven adverse determination or clinical decision takes effect.
CMS Conditions of Participation
CMS expects hospitals to maintain clinical judgment standards. An algorithm cannot satisfy the physician's independent duty of care regardless of its accuracy rate.
What 'review' must actually mean
Not a rubber stamp
Clicking 'approve' on every AI output in under 3 seconds does not constitute review. Courts and auditors will examine click-through rates and time-on-task logs.
Documented reasoning
When the reviewer agrees with AI output, a brief note suffices. When they override, the reasoning must be captured — this is your FCA defense document.
Clinical authority must be real
The reviewer must have the authority and training to override the algorithm. Designating a reviewer who lacks the credentials to challenge AI output is not compliance.
Audit trail integrity
Every review, approval, and override must be timestamped and immutable. The absence of an audit trail is itself evidence of reckless disregard under the FCA.
The UCHealth lesson: The algorithm triggered the overpayment. The absent human review triggered the $23M liability.

13

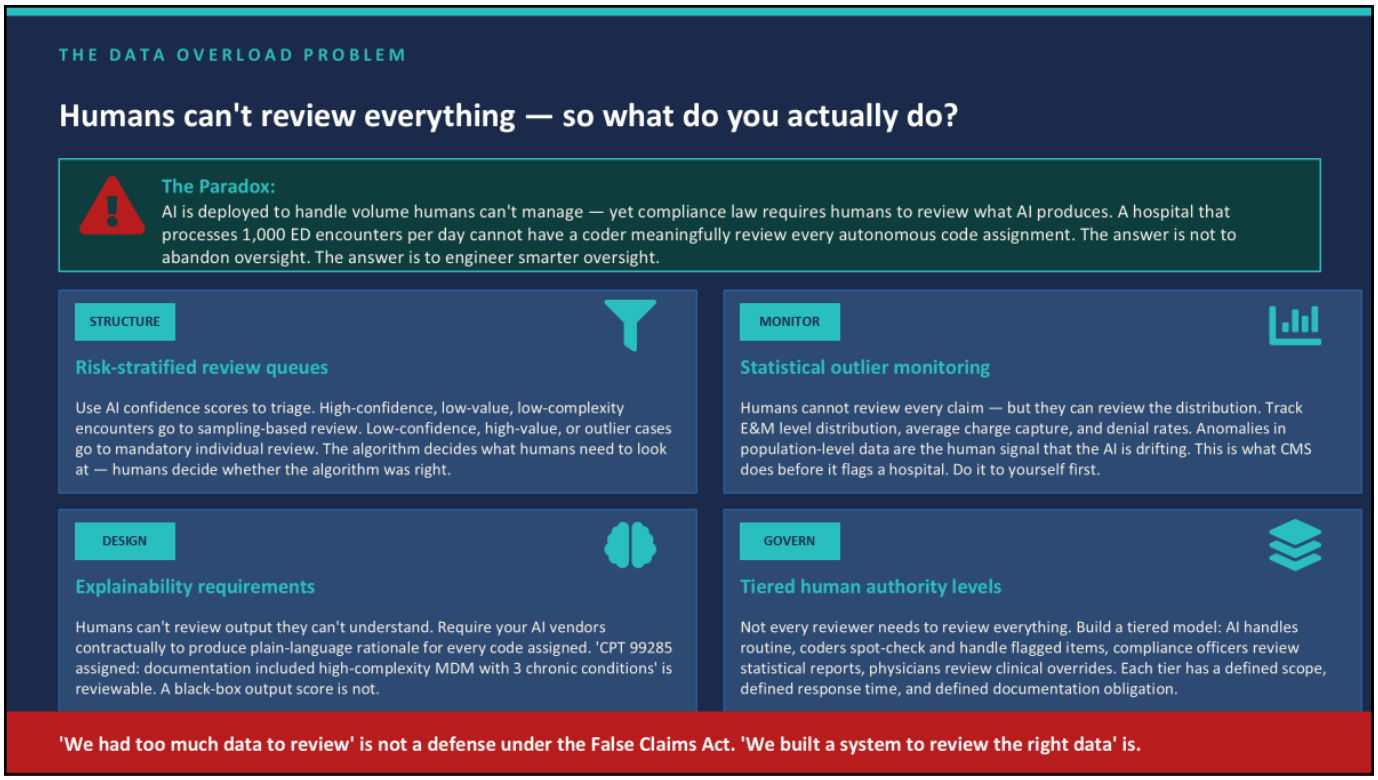
THE DATA OVERLOAD PROBLEM
Humans can't review everything — so what do you actually do?
The Paradox:
AI is deployed to handle volume humans can't manage — yet compliance law requires humans to review what AI produces. A hospital that processes 1,000 ED encounters per day cannot have a coder meaningfully review every autonomous code assignment. The answer is not to abandon oversight. The answer is to engineer smarter oversight.
STRUCTURE
Risk-stratified review queues
Use AI confidence scores to triage. High-confidence, low-value, low-complexity encounters go to sampling-based review. Low-confidence, high-value, or outlier cases go to mandatory individual review. The algorithm decides what humans need to look at — humans decide whether the algorithm was right.
MONITOR
Statistical outlier monitoring
Humans cannot review every claim — but they can review the distribution. Track E&M level distribution, average charge capture, and denial rates. Anomalies in population-level data are the human signal that the AI is drifting. This is what CMS does before it flags a hospital. Do it to yourself first.
DESIGN
Explainability requirements
Humans can't review output they can't understand. Require your AI vendors contractually to produce plain-language rationale for every code assigned. 'CPT 99285 assigned: documentation included high-complexity MDM with 3 chronic conditions' is reviewable. A black-box output score is not.
GOVERN
Tiered human authority levels
Not every reviewer needs to review everything. Build a tiered model: AI handles routine, coders spot-check and handle flagged items, compliance officers review statistical reports, physicians review clinical overrides. Each tier has a defined scope, defined response time, and defined documentation obligation.
'We had too much data to review' is not a defense under the False Claims Act. 'We built a system to review the right data' is.

14

CASE STUDY 03

Banner Health | Phoenix, Arizona (AZ / CA / CO)

AI Implementation

Banner Health deployed a broad RPA and machine learning strategy across its revenue cycle, creating bots for insurance coverage discovery, claims correspondence, and — most significantly — a predictive model to determine when a write-off is warranted based on denial codes and probability-of-payment analysis.

False Claims Act Risk

Automated write-offs bypass human judgment. If a bot incorrectly writes off a collectible claim, this may constitute improper billing or failure to collect required co-pays — both FCA exposure areas.

Model Validation Requirements

CMS and OIG guidance increasingly requires that AI/ML models used in billing be clinically and financially validated before deployment. Banner must document model accuracy and audit trails.

The Danger of Policy Changes or Criteria Updates

Healthcare is a constantly changing environment. Coding guidelines and even the codes themselves are often updated. AI solutions must be able to change with them, and oversite must be in place to ensure those updates have been added to the models.

"Machine learning and predictive analytics are great tools... Soon, AI will synthesize that data and make improvement recommendations." — Banner Health RCM Automation Director

15

WHO PAYS WHEN AI GETS IT WRONG?

Hospital Liability, AI Vendors & the False Claims Act

The law holds the entity that submits the claim responsible — not the vendor whose algorithm generated it.

False Claims Act

31 U.S.C. § 3729

Liability attaches to the claim submitter. 'Knowingly' includes reckless disregard — not just intent. Every AI-generated code submitted to Medicare is your legal exposure.

HIPAA Covered Entity

45 C.F.R. Parts 160 & 164

The hospital is the covered entity. Every AI vendor touching PHI must have a signed BAA. Using AI without one is a violation regardless of clinical intent.

Respondeat Superior

Common Law Doctrine

Hospitals are vicariously liable for how employees use AI within the scope of work. Corporate negligence adds an independent duty to vet and monitor every AI tool deployed.

Citations: FCA 31 U.S.C. §§ 3729-3733 | HIPAA 45 C.F.R. Parts 160 & 164 | AKS 42 U.S.C. § 1320a-7b | Revenue Cycle Conference

16

REAL-WORLD PROOF

UCHealth · $23M FCA Settlement

$23M

DOJ Settlement
Nov. 2024

The Algorithm:
AI auto-coded all ED visits as CPT 99285 (highest level) based on vital sign checks — regardless of actual acuity.

The Red Flag:
CMS flagged UCHealth as a 'High Outlier.' Coding staff complained repeatedly. Hospital did not adjust the rule.

DOJ's Theory:
'Reckless disregard' — not fraud. No intent required. The vendor wrote the algorithm; UCHealth paid the $23 million.

WHAT YOU MUST DO

Six Non-Negotiable Actions

01 Audit AI outputs regularly — if you can't explain the algorithm, you can't defend it to DOJ

02 Monitor E&M distribution; 'High Outlier' flags from CMS are effective DOJ notices

03 Execute BAAs with every AI vendor touching PHI — no exceptions

04 Document when clinicians override AI recommendations and why

05 Never ignore internal coding complaints — they become whistleblower evidence

06 Update your compliance program specifically for AI; generic programs leave FCA gaps

Automated systems are NOT a compliance shield. The vendor writes the algorithm. You sign the Medicare claim.

17

CASE STUDY 04

Baptist Health | Little Rock, Arkansas

Powered by
AccuCodeAi

Automating data abstraction for clinical quality measures in value-based care where CQM performance directly determines Medicare and managed care reimbursement.

THE SITUATION

Baptist Health System

One of the largest not-for-profit health systems in Arkansas, serving patients across a broad network of hospitals and outpatient facilities in central and northwest Arkansas.

Under value-based care contracts, clinical quality measure (CQM) performance directly determines reimbursement rates. Manual abstraction of the data required to report those measures was time-intensive and consume significant clinical and administrative staff resources. AccuCode's ability to review every piece of clinical documentation unlocked new insights to improve patient care.

AI APPROACH

Automated CQM abstraction

AccuCode AI's engine reads clinical documentation and automatically abstracts data elements needed for quality measure reporting replacing manual chart review at scale.

Humans kept in the loop

Baptist maintained credentialed clinical reviewers at every validation checkpoint. AI outputs are reviewed before submission, a governance decision that has been central to the program's success and defensibility.

Expanded opportunity discovery

The AI surface-mining process identified additional improvement opportunities across Baptist's operations that manual abstraction had never surfaced, creating value well beyond the original project scope.

COMPLIANCE CHALLENGES

Sheer volume of returned data

Abstraction surfaces far more data than manual workflows, including elements requiring clinical judgment to interpret. Managing that volume without overwhelming reviewers, or losing signal in the noise, is an ongoing operational challenge.

Humans in the Loop Own the Submissions

Baptist Health retains full ownership of every CQM submission to CMS. AccuCode AI's workflow ensures human-in-the-loop validation at every stage, keeping Baptist's quality leaders in control of the final record.

Value-based payment exposure cuts both ways

Inaccurate abstraction can inflate or deflate performance scores. Overstated scores invite CMS audit and potential repayment demands. Understated scores forfeit earned reimbursement. Both directions carry real financial and compliance risk.

Key success factor: Baptist kept humans in the loop, not as a checkbox, but as a genuine governance decision. That choice is the difference between a defensible AI deployment and a liability event.

CQM Automated quality measure abstraction

VBC Value-based payment at stake

HitL Human-in-the-loop validation model

18

Abstraction Accuracy Performance

Baseline to Superhuman to Best-in-Industry

>95%
Superhuman Accuracy
On First Pass

830+
Unique Data Elements
Built and Tested

1,170+
Total Element Instances
Across All Measures

The Refinement Trajectory

3	11	14	13
Best-in-Industry	**Superhuman**	**Built**	**New Additions**
Confirmed	Confirming refinement	Awaiting refinement	Awaiting patient samples

AccuCode's initial output, before any human review, consistently exceeds CMS compliance targets. In every instance, the overwhelming majority of discrepancies were resolved in AccuCode's favor.

19

Clinical Discoveries & Patient Care Impact

AccuCode's comprehensive record analysis across all our clients identified issues in care processes that were not previously visible.

BP Monitoring Correction

Discovered that blood pressure readings were recorded at incorrect times due to an equipment connectivity issue. Devices batched and transmitted readings upon network reconnection rather than at time of measurement. Clinically significant events were being logged at the wrong time, potentially delaying response. Corrected as a direct result of AccuCode's analysis.

Expansion of the Accessible Medical Record

Because AccuCode reads the complete medical record, which can be an many as 175 pages of single-spaced text, it has the potential to discover early documentation of important events. Every medical record in the patient population can be abstracted and returned in a matter of hours for an entire quarter.

Potential for Documentation Improvements

- **Crystalloid Fluid Administration** Incomplete volume documentation for sepsis patients. Nursing workflows being updated.
- **Sequential Compression Devices** Devices ordered but lacking documented confirmation of application. Standards being revised.
- **Patient Education Delivery** Attestation step (nurse sign-off) frequently missing. Being added to discharge workflows.
- **Physician Sepsis Documentation** Physicians being coached to align with guideline-specific language requirements.

20

Cross-Case Compliance Themes

What Every Case Study Has in Common

1. AI Errors = Hospital Liability

Across all three initial cases, when AI produces a billing error — whether a miscoded claim, an incorrect write-off, or a misclassified level of care — the hospital bears full legal and financial responsibility under the False Claims Act and CMS rules. Vendors are not liable.

2. Human Oversight is Non-Negotiable

No case study succeeded with fully autonomous AI. Auburn used coders to review AI suggestions. Inova built new governance teams. Northwell required human chart reviewers. OIG and CMS both expect human review checkpoints in AI-assisted billing.

3. HIPAA Compliance Starts Before Go-Live

Every AI vendor touching PHI must be vetted as a Business Associate before any data flows. Several organizations underestimated the infrastructure cost of maintaining HIPAA-compliant AI environments — especially for LLM fine-tuning.

4. Regulatorily Evolving Standards

Coding standards (ICD-10, CPT updates), CMS payment policies, and RAC audit priorities change continuously. CMQ measure are updated, dropped or added every year. AI models trained on historical data can become non-compliant without regular retraining and compliance auditing.

21

AI in Revenue Cycle Management

5-Year Market Projections

$20.6B
Market size 2024

$70.1B
Projected 2030

~24%
CAGR 2025-2030

67%
Providers using AI

Sources: Grand View Research · McKinsey · HFMA · Experian Health 2025

AI RCM Market Growth 2024 - 2030

AI-specific revenue cycle management market size (USD billions)

42%	70%	30%	98%
Denial rate reduction	Coding error reduction	Admin burden cut	Providers plan AI by 2026

22

The Opportunity is Real.

The Compliance Stakes Are Higher.

AI can transform your revenue cycle — but only with the right governance, human oversight, and compliance infrastructure in place.
The hospitals that win will be those that treat compliance not as a constraint on AI, but as the foundation for it.

Scott Roper, VP Business Development, AccuCodeAI, Inc.
Nathan Myers, CEO AccuCodeAI, Inc.

AR HFMA | 2026

23

INTERACTIVE SESSION

Your Turn: AI in the Revenue Cycle — Share, Reflect & Envision

We've explored five real-world case studies. Now let's hear from the room. For each topic below, share your experience — what worked, what didn't, and what you want next.

AI Successes *What worked — and why*		AI Failures *What didn't work — and what you learned*		Future Vision *What you want AI to do next*	
What AI project delivered your best ROI?	Wins	Where did an AI implementation fall short or fail?	Failure	What RCM problem do you most want AI to solve?	Priority
Which compliance risk did AI help you avoid?	Compliance	What compliance issue did AI unexpectedly create?	Risk	What compliance capability is missing from current AI?	Gap
What surprised you most about AI performance?	Insight	How did staff resistance derail a project?	Culture	Where do you see AI in your RCM in 3 years?	Strategy
How did your team adopt AI most successfully?	Culture	What did a vendor promise that wasn't delivered?	Vendor	What would make you more confident deploying AI?	Trust
Which vendor or tool exceeded expectations?	Vendor	How did a model perform well in testing but fail live?	QA Gap	What governance structure do you wish existed?	Policy

24