CARTA HEALTHCARE

WHY HYBRID INTELLIGENCE

Is the End State for Clinical Data Abstraction

A Strategic Perspective on AI, Governance, and the Future of Registry Abstraction

In an era of rapid AI advancement, the question of whether human oversight remains necessary in clinical workflows deserves a rigorous answer. This paper makes the case that hybrid intelligence — not full automation — is the durable architecture for regulated healthcare.

Carta Healthcare · 2025 · Whitepaper

A Question Worth Answering

Frequently, we get a fair question which is: as AI continues to improve, won't horizontal, AI-only companies eventually overtake us? It is the kind of question you would expect astute people to ask. AI models are advancing rapidly. Automation rates are climbing. The cost of intelligence is falling. If clinical abstraction is fundamentally about reading charts and extracting data, it feels reasonable to assume that full automation will eventually win.

That logic holds if abstraction is simply a model capability race. But it is not. Clinical data abstraction operates inside a regulated, audited, high-consequence system. Registry data does not live in isolation. It feeds CMS reimbursement, accreditation status, public reporting, quality programs, performance dashboards, and population health strategy. When the data is wrong, it is not just a technical miss — it can create audit exposure, financial consequences, and reputational risk. That reality changes the optimization function entirely.

> ***The architecture that wins in an exception-heavy environment is the one designed around exception handling and accountability — not just throughput.***

CROSS-INDUSTRY EVIDENCE

Automation With Accountability: A Universal Pattern

When you look across other high-stakes industries, a consistent pattern emerges. In each case, automation increased dramatically — but oversight remained embedded in the architecture. Clinical abstraction belongs in this same category.

Aviation

Modern aircraft can largely fly themselves, and autopilot systems are extraordinarily precise. Yet we did not remove pilots from the cockpit. Aviation is not optimized for maximum automation; it is optimized for risk-adjusted performance. Redundancy, escalation capability, and accountability remain embedded because the question is not whether the system can *usually* fly the plane. The question is what happens when something unusual occurs.

Radiology

AI models can detect nodules and flag abnormalities with impressive sensitivity. Yet radiologists were not eliminated as the technology improved. Instead, the model evolved into AI as a second reader — and clinicians retain final interpretation and legal responsibility. The legal and regulatory frameworks did not disappear simply because the models improved. The same applies to abstraction: hospitals will not remove oversight because a model improves incrementally. They will continue to ask who validates the data and who is accountable when an auditor shows up.

Finance

Algorithmic trading dominates execution speed, but when volatility events exposed systemic risk, the industry responded by adding governance — not removing it. Kill switches, monitoring desks, compliance layers, and circuit breakers became part of the architecture. Speed without guardrails proved dangerous. Registry abstraction also carries systemic implications: errors cascade into dashboards, public metrics, accreditation reviews, and reimbursement. Institutions will always demand governance where downside risk is asymmetric.

AI-Driven Customer Support

Early on, the assumption was that chatbots would replace support agents entirely. For simple tasks like password resets or order status checks, they did. But when issues became nuanced, emotionally charged, or financially sensitive, companies learned that full automation eroded trust. The systems that worked best were not fully autonomous — they used AI to handle routine inquiries at scale while designing seamless escalation to humans for complex cases. Healthcare abstraction operates under far higher stakes, and the winning architecture follows the same pattern.

CORE THESIS

Hybrid Intelligence as the Structural End State

This is why we do not view hybrid intelligence as a temporary bridge to full automation. I see it as the structural end state for regulated healthcare workflows. Hybrid is not simply AI plus humans layered together — it is a control architecture.

How the Architecture Divides Responsibility

AI handles: pattern recognition at scale, temporal synthesis across fragmented documentation, and routine case throughput.

Human experts handle: ambiguity, contextual interpretation, edge cases, and final accountability.

The system is designed around escalation and defensibility rather than pure automation. Clinical abstraction is full of unusual situations — incomplete documentation, ambiguous timestamps, conflicting records, registry rule nuances, and subtle comorbidities buried in unstructured notes are common, not rare. AI-only systems tend to optimize for the routine case. Healthcare, however, is an exception-heavy environment.

THE CARTA HEALTHCARE DIFFERENCE

Our Objective Function Is Different

The difference between Carta Healthcare and horizontal AI-only companies is not simply model quality — it is an objective function. Horizontal companies optimize for automation rate, scalability, and margin. We optimize for sustained 98–99% accuracy, audit readiness, registry-specific nuance, governance workflows, and institutional trust.

Even if an AI-only company eventually matches raw extraction performance, they still have to solve for embedded oversight, IRR management, adjudication processes, and integration into the operational fabric of health systems. They can replicate intelligence. It is much harder to replicate a risk-governed operating system.

> ***The core insight: AI capability and institutional risk tolerance do not move at the same speed. That gap is durable. Hybrid intelligence lives inside that gap.***

THE REAL STRATEGIC RISK

The System Is the Moat

The real strategic risk for us is not that AI improves — it will. The risk would emerge if abstraction were reframed internally within health systems as purely a technology function and governance became something they were comfortable owning entirely themselves. That is why our long-term positioning matters. We cannot simply be an AI vendor. We need to be the operating layer for risk-governed clinical abstraction.

If we own the workflow, the governance infrastructure, the registry expertise, and the audit interface, then the model underneath can evolve indefinitely. The model is not the moat. The system is.

CARTA HEALTHCARE

CONCLUSION

Hybrid Intelligence Is Not a Compromise. It Is the Architecture.

When you step back and look across aviation, radiology, finance, and other high-risk industries, the pattern is consistent: automation increases, but accountability remains embedded. The dominant architecture becomes automation with oversight.

Clinical data abstraction sits squarely in that category. Hybrid intelligence is not a compromise between humans and machines. It is the architecture that aligns automation with governance in a high-consequence environment. That is why it is the durable end state — not a transitional phase.

Strategic Questions

Q: As AI improves, could AI-only competitors overtake us?

A: That assumes abstraction is a model race. It isn't. It's a regulated workflow tied to reimbursement and accreditation. In high-consequence systems, the durable architecture is automation with accountability, not automation alone.

Q: If models become near-perfect, won't human oversight become unnecessary?

A: Institutional risk tolerance does not rise at the same speed as model capability. In aviation, finance, and radiology, automation increased, but oversight remained embedded. Healthcare follows the same pattern.

Q: What truly makes Carta Healthcare different?

A: The solution is the system, not the model. Our differentiation is registry expertise, governance workflows, adjudication infrastructure, IRR management, audit defensibility, and operational integration. Those are difficult to replicate.

Q: What is the real long-term risk to the business?

A: The risk is not AI improvement. The risk would be if abstraction is reframed as purely a technology function and hospitals internalize governance. Our strategy mitigates that by positioning Carta as the operating layer for risk-governed abstraction.

Q: Is hybrid transitional or permanent?

A: In regulated industries with asymmetric downside risk, the dominant architecture is automation with embedded oversight. Hybrid is not a bridge. It is the structural end state.

Carta Healthcare · Where AI-driven innovation meets real-world healthcare challenges.

carta.healthcare