Within3

E-BOOK



An AI reckoning:

Why Life Sciences Has More AI Than It Can Use *and What's Still Missing: The Intelligence Layer*

Table of Contents

Introduction 3

Part 1: The Reckoning 4

Part 2: The Reasons 5

- Workflow automation ≠ intelligence 6
- Bigger ≠ better, and generalization ≠ innovation 7
- Governance as afterthought ≠ agility 9
- Volume ≠ value 11

Part 3: The Response 12

- What is the intelligence layer? 13

Conclusion 15

Within3

With any new technology, there comes an inevitable saturation point. Has life sciences reached 'peak AI'? The first internet-enabled BlackBerry cell phone was released in 2002, while the Apple iPhone was announced in January 2007, both major milestones in smartphone technology. Barely 20 years on, there are now 7.5 billion smartphones in active use worldwide. Looking a little further back, the first 'website' worthy of the name went live in 1991, but by 2014, there were more than one billion of them. AI deployment hasn't quite reached that stage, but it's getting there quickly. In 2025, 1,953 new AI companies were founded in the US alone, a jump of 71% from the previous year. In life sciences specifically, 79% of companies now report active AI use.

When a technology reaches a saturation point, several things happen:

- User confidence and familiarity increase
- Adoption costs come down
- Growth flattens
- Competition intensifies
- Differentiation collapses
- Consolidation begins

> "The market is flooded with me-too models. As one analyst noted, 'the race isn't just about who builds the best algorithm – it's about who can monetize it first'." – *AI INVEST*

Now, we're beginning to see these trends in pharma AI. The result for the end user is that it becomes more and more difficult to separate the wheat from the chaff. The more potential outlets you have for your AI investment, the harder it becomes to put those dollars to effective use. The time has come for a pharma AI reckoning. The market is established. The technology is mature. But are AI deployments really delivering material value in terms of brand performance and product adoption in the context of a highly regulated pharmaceutical industry? Or is it just individual productivity gains? If you're a pharma leader and your answer is not aligned with the former, you're leaving a lot of promise on the table.

PART 1: The Reckoning

> "The North America AI in pharma & biotech market size was valued at USD 2.83 billion in 2024. The market is projected to be worth USD 3.43 billion in 2025 and reach USD 39.78 billion by 2032."
>
> – *FORTUNE BUSINESS INSIGHTS*

Life sciences AI investment is at an all-time high, and as these figures show, it's growing exponentially. Deployment announcements are near-constant, and every major vendor is shipping AI capabilities, embedding agents, and expanding its footprint across commercial and medical workflows. Comparatively speaking, this is still a new market: one that didn't exist a decade ago, but within the next five years, projections place it at almost $40 billion in the US alone. We're fast approaching an AI saturation point, if we haven't already passed it. And yet the question remains: are the majority of these deployments actually delivering measurable business impact? Ask the organizations responsible for them – in private – and they might admit to an uncomfortable truth.

McKinsey reports that, as of 2025, 88% of organizations are using AI in at least one function, with 64% agreeing that AI enables innovation. Just 39% of respondents, however, reported EBIT impact at the enterprise level as a result of their AI investments. An even smaller proportion, just 6%, could report an EBIT impact of ≥5%. Seemingly, life sciences AI adopters find themselves in the role of Coleridge's Ancient Mariner, with "water, water, everywhere, nor any drop to drink." They're drowning in a sea of AI tools, but the value of those deployments remains debatable at best.

Put bluntly: an AI reckoning is due. Most pharma organizations have passed the 'why invest in AI' phase and have reached the 'where' and the 'how'. They're spending huge amounts of money on AI deployments, around $4 billion by 2025, and could reasonably expect a significant return on that investment. If vendors can't begin to demonstrate compounding business impact from all of those deployments, pharma will begin asking the 'why AI' question once more. But this isn't a technology problem – the models are capable, the infrastructure exists, and vendors ship promptly. Rather, a gap exists between what AI is producing and what pharmaceutical organizations actually need in order to execute high-stakes decisions. Until that gap is closed, more deployments will not produce more impact, only more activity.

Fortunately, there are several clear reasons why pharma AI deployments aren't currently delivering on their early promise, as well as a single clear solution to those problems.

PART 2: The Reasons

AI deployments are failing to deliver tangible value to pharma organizations because some of the industry's biggest vendors have proposed the wrong solutions to pharma's challenges. It's less a technological failing than one of market positioning and utilization. If pharma has been presented with a bowl of soup, AI vendors have so far provided only forks. While the tool is perfectly functional, there are issues in how it's being put to work. Consider the following structural failures:

1 WORKFLOW AUTOMATION ≠ INTELLIGENCE

If pharma AI deployments are proliferating, and those deployments largely aren't delivering tangible value, there's a certain logic to the idea of consolidating AI investments in one place. CRM vendors such as Veeva and Salesforce are beginning to offer AI agents within their own ecosystems, allowing teams to automate essential workflows without the need for further external AI investment. CRM systems and content tools have long been designed to capture and manage activity: calls logged, meetings scheduled, content distributed. Layering AI into these systems enables automation of several essential processes. There's real value in this approach, and the pitch is certainly appealing: a single platform, a single contract, a single vendor relationship. This is the 'suite as hub' argument, eschewing extraneous, expensive, underperforming AI investments in favor of highly functional integrated AI agents at a minimal outlay.

The logic may be seductive, but while this approach applies to workflow tools, it is not applicable to intelligence infrastructure. This is the first issue with current pharma AI deployments: the industry has conflated workflow automation with genuine intelligence. In reality, launch decisions are not made within any single workflow system but by synthesizing signals across many: field activity, advisory boards, claims data, real-world evidence, HCP sentiment, competitive intelligence, and more. Further, attempting to combine workflows and intelligence into a single suite undercuts the capabilities of both. AI embedded within a workflow system produces little more than an efficient view of a partial reality. It organizes what the system can see, but doesn't account for what it cannot. As Within3's Chief Technology Officer, Jason Smith, has argued across a recent series on these structural gaps:

> "A CRM optimized for intelligence production becomes a worse CRM. An intelligence layer built inside any single workflow platform cannot be truly neutral across the multiple enterprise systems it needs to integrate." – *JASON SMITH, CTO WITHIN3*

Despite how the industry is currently positioning itself, the optimum AI architecture is additive, not consolidated.

Learn why market truth lives above your CLM

Read blog ↗

Within3

2 BIGGER ≠ BETTER, AND GENERALIZATION ≠ INNOVATION

There's a similar industry misconception around model size in pharma AI. The AI arms race has led vendors to produce larger models with broader capabilities and parameters. The seemingly logical argument is that 'bigger is better', that capability scales in proportion to model size, and that the largest general-purpose LLM will be best positioned to provide any and all answers — a silver bullet. While it's true that general-purpose LLMs are great at providing convincing, well-written responses to teams' queries, it's less certain whether those persuasive responses can be trusted, especially for pharma.

Up to now, the industry has been passing off generic AI tools as 'innovation'. However, medical affairs engagements, competitive positioning, or market access barriers are anything but generic, and general-purpose tools lack the context and specificity to extract real insight from such data. For that reason, general-purpose LLMs should never be equated with purpose-built pharmaceutical intelligence, and AI-assisted summarization tools packaged within general workflow systems should be viewed as insufficient for medical insights capture.

The truly problematic issue is that these general-purpose LLMs are good at creating convincing, plausible summaries. If they weren't, errors and gaps would be easier to identify and rectify. With a general-purpose tool, teams run the risk of making critical launch decisions based on hallucinations or incomplete information. General-purpose LLMs aren't providing teams with the truth, but a convincing simulacrum of it. This distinction between plausibility and true fidelity, while small in casual use cases, is highly consequential in fields such as pharma.

> "LLMs have enormous transformative potential. However, LLMs also present challenges, e.g. variability in results, returning of irrelevant or inaccurate responses (so-called hallucinations), and potential data security risks."
>
> – *EUROPEAN MEDICINES AGENCY*

AI performance can only be evaluated in the context of use. As Jason Smith, CTO of Within3, has argued: “The output quality question is not ‘Is this well-written?’ It is ‘Is this true?’”. Pharma doesn’t need the biggest general-purpose LLM, but the most specific and contextual: a small language model trained on industry data and capable of offering the closest ‘domain fit’. The difference between the two models in terms of output is subtle, but critical: it’s the difference between ‘contextually plausible’ and ‘factually accurate relative to the source material’. No matter how large the model, a general-purpose LLM will never deliver true value in a highly specialized, highly regulated industry like pharma, where assumptions alone can’t provide a solid foundation for decision support.



“The differentiation that matters in this space is not ‘more content, faster.’ It is ‘fewer wrong decisions, with evidence.’ That positioning requires a credible evaluation program, benchmarks against domain-specific task sets, and transparency to be honest about what a model does not do well, alongside what it does well. That transparency, backed by evidence, is what earns trust in a regulated environment.” *– JASON SMITH, CTO WITHIN3*

Learn the five capabilities of a domain-fit AI model

Read blog ↗

Within3



3 GOVERNANCE AS AFTERTHOUGHT ≠ AGILITY

Within3's 2026 AI predictions report revealed some fascinating insights into how industry leaders view pharma's regulatory landscape. While 42% of survey respondents expected stricter privacy and compliance requirements to slow innovation in 2026, 19% thought that clearer guardrails would actually accelerate innovation by increasing confidence and adoption. Regardless of your position, it's clear that compliance is becoming a competitive differentiator, with those capable of building smart governance processes around data privacy and security, while accepting reasonable risk, likely to outperform those that aren't.

> "Compliance is going to become a competitive issue. If you can't adopt quickly, you can't innovate – and you'll miss the opportunity."
>
> *– TONY PAGE, SENIOR VICE-PRESIDENT, INSIGHT ANALYTICS*

The problem for pharma is that the AI industry has built governance as an afterthought and packaged it as 'agility'. Today, every major life sciences AI vendor will claim that governance is 'built in' – effectively, that you needn't worry, because the product has it covered. That's their case for 'agility'. But when governance is implicit rather than explicit, every new workflow requires fresh compliance justification. Scaling any new AI deployment costs just as much time, resources, and political capital as the very first time. That's the opposite of true agility.

> "Governance needs to be demonstrable, not implied. A compliance committee cannot approve what it cannot examine."
>
> *– JASON SMITH, CTO WITHIN3*

Similarly, attempting to bypass governance through open-ended autonomy means that outputs cannot be traced to their source data, systematically validated, or audited retroactively. With this approach, risks are liable to emerge only at the most inopportune times: during regulatory scrutiny, for example, or following commercial underperformance. A launch signal that proves to be inaccurate and cannot be traced to its source will create legal, regulatory, and commercial exposure that will far exceed the cost of baking governance into the architecture from the beginning.

The organizations that will scale AI moving forward are those treating governance as a product feature rather than a compliance tax. As Jason has argued, "The AI evidence pack is not documentation overhead. It is a commercial accelerator." When that pack exists – model version, datasets, evaluation results, limitations, and human sign-off documented per workflow – a compliance conversation becomes a handoff. When it doesn't, it becomes an investigation.

Learn how to scale AI effectively in 2026 and beyond

Read blog ↗

Within3

4 VOLUME ≠ VALUE

Continuing the pattern of industry mislabeling, pharma AI has conflated output with impact. But for pharma organizations, the limiting factor has never been an absence of information, but rather the ability to translate information into decisions at speed. For many organizations, decision velocity still lags. A general-purpose generative AI tool, deployed without structural constraints, will inundate launch teams with summaries, reports, and synthesized content. But unless those outputs are aligned with launch strategy and routed to the correct decision-makers, they become a thicket of information that must be cut through before any decision is reached. Far from accelerating decision velocity, more AI output equates to more noise for leaders to navigate.

Currently, pharma's AI problem is not a lack of output but a lack of meaningful, decision-relevant signal. 'More volume' simply contributes to information overload in overwhelmed decision-makers. To deliver real value, pharma AI has to improve its 'signal-to-noise ratio' and deliver outputs that actually affect strategic decision-making. The goal must be 'strategy-aligned AI': a system designed to address a small number of high-stakes strategic questions, not to generate the most summaries.

> "The difference between an AI deployment that scales and one that stalls is almost never the model. It is whether the outputs are connected to the strategic questions that decision-makers are actually trying to answer." – *JASON SMITH, CTO WITHIN3*

Further complicating the issue is the fact that current AI tools excel at data aggregation. The industry calls this 'completeness', when in reality the picture is fragmentary. Without governance that travels with the data – identity resolution, quality certification, privacy controls, auditability – data aggregation produces a facsimile of comprehensiveness, but without trustworthiness. A signal without trust isn't actionable.

PART 3: The Response

We started this article by asking whether life sciences had reached 'peak AI', and it's easy to argue that it has. That would be the provocative headline. But as the issues outlined above demonstrate, it's not that there's too much AI, just that it's being deployed in ineffective or inefficient ways to truly drive meaning for pharma teams. Teams needn't walk back their AI investments per se, but they may want to rethink where and how that money is allocated.

In part 2, we discussed how the optimum architecture for AI deployment is additive, not consolidated. The wider industry push toward consolidation, where AI is being bolted onto existing systems, is misguided. But what might an additive architecture look like? Successful AI deployment should be thought of as an architecture where the intelligence layer wraps around all other foundational layers in the ecosystem.

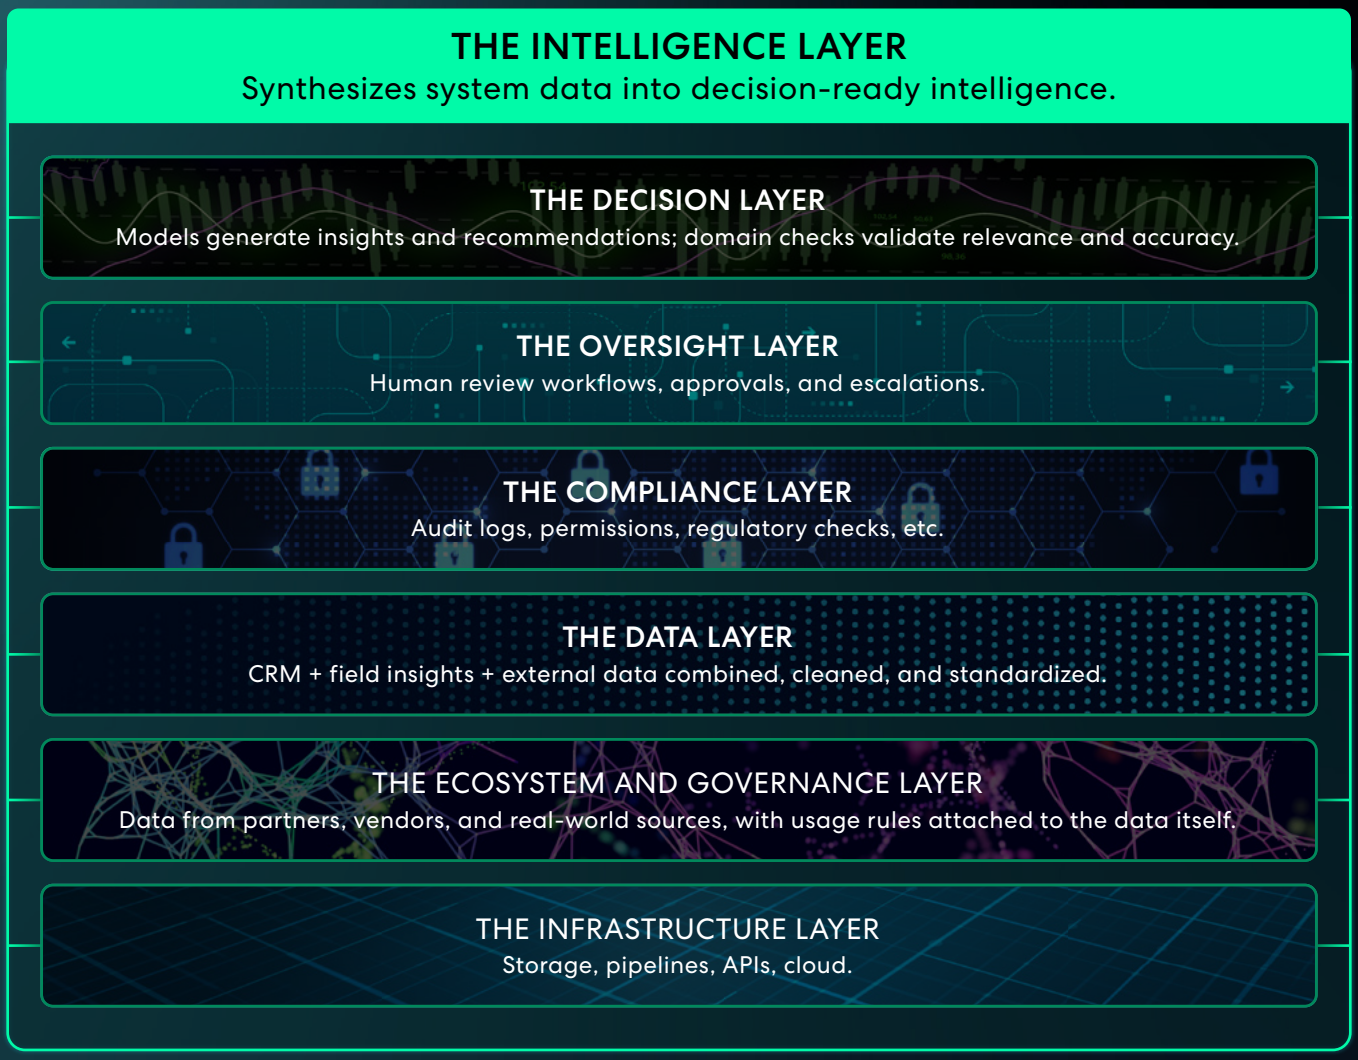

In 2026, most pharma companies have some of these layers in place already, although precious few have all of them working in harmony. Without the full stack, AI can produce outputs, but not compound them into business impact. And consistently, the 'missing layer' is the final, most valuable one: the intelligence layer.

WHAT IS THE INTELLIGENCE LAYER?

What's missing is not another model, dashboard, or workflow tool. It's the intelligence layer that sits above existing systems and does the work those systems were never designed for.

> "An AI system that processes only the signals it owns natively is not producing launch intelligence. It is producing a well-organized view of a subset of the available evidence, dressed up as a complete picture."

– JASON SMITH, CTO WITHIN3

Before now, the industry invested in systems that excel at capturing workflow activity within their own boundaries. These have value in isolation, but they are not holistic. Where the industry has underinvested is in an additional layer capable of synthesizing signals across those boundaries and producing outputs that decision-makers can actually act on – what can be thought of as 'decision-ready intelligence'. That's the intelligence layer, and it's the gap where AI transforms from a costly investment to one capable of delivering real business impact.



This layer has several discrete components:

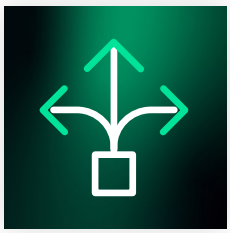

Signal unification above workflows

True intelligence requires the ability to synthesize across systems without bias. That means resolving identities across datasets and integrating signals in a way that preserves context and comparability.

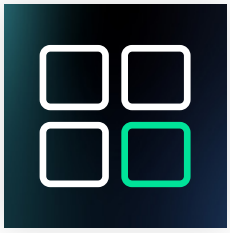

Domain fit evaluation

AI systems can only be tested against the tasks they were asked to perform, using real data and defined standards. The metric that really matters is not some abstract performance benchmark, but how accurate outputs are when checked against the source material.



Compliance-native operations

Governance must be embedded into the architecture itself: audit logs, role-based access controls, policy-driven output constraints, and traceable provenance all the way from data to decision.



Human oversight as a designed control

'Human-in-the-loop' is the only mechanism through which organizations can generate the evidence required to trust and scale AI deployments. Structured review processes, feedback loops, and drift monitoring transform human validation into a source of continuous system improvement.

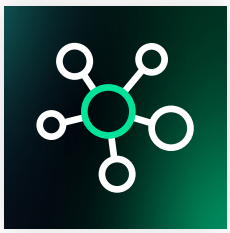

Ecosystem data coverage with embedded governance

The intelligence layer must integrate external data sources (claims, EHR, real-world evidence, expert engagement) with the governance structures that make them usable (privacy-preserving linkage, data quality certification, and compliance templates that travel with each source).

The pharma companies behind the most successful launches of the next half-decade will not be defined by the CRM they use, nor the AI agents embedded within that CRM. What will set them apart is their ability to leverage a dependable, fully governed intelligence layer that sits above their systems and delivers decision-ready intelligence at the speed of launch.

In conclusion

We are, collectively, still learning about AI. We're learning what it's good at, what it isn't, and more importantly, how to use it. Most critical thinkers have already recognized that general-purpose LLMs have limitations and drawbacks, but that doesn't mean they've stopped using them altogether. They've simply adapted, using AI agents in different, more creative, and more effective ways.

In the same way, a pharma AI reckoning may have come due, but the industry would be wrong to throw the baby out with the bathwater. Instead, the time has come to ask why AI deployments aren't delivering tangible value and how organizations can rethink their investments and adjust their approach to change that dynamic. Jason put it plainly: "The fastest AI in a regulated organization is not the one that generates the most content. It is the one that generates content that clears review without revision on the first pass."

AI isn't an illusion; it's not some flash in the pan, and it will revolutionize pharmaceutical launches. However, the pharma companies that succeed in the years to come will not deploy the most AI, but the AI that matters most.

AI platforms designed around an additive architecture with the intelligence layer at the top, Within3's Launch Intelligence™ among them, are directionally aligned with what the industry actually needs. Not more output, but better decisions, made at pace, and verifiable through demonstrable governance. In this way and this way alone, AI will finally deliver on its great promise.

ABOUT WITHIN3

Within3 empowers pharma leaders to make confident and timely launch decisions. Our Launch Intelligence™ platform unifies field insights, stakeholder engagement, social, claims, and market signals into a single connected view that surfaces what matters most. Powered by a life-sciences trained AI architecture and guided by expert strategists, Within3 helps launch teams move with clarity and speed, turning complex evidence into decisive action.

Trusted by all of the top 20 pharmaceutical companies, Within3 is the partner modern launch leaders rely on to align teams, act faster, and accelerate outcomes across the product lifecycle.

Within3