Article

AI and Analytics Evolution within Commercial Life Sciences

CHRISTIE CONN, VP and GM, Analytics & Technology, IQVIA
GOKSU DOGAN, Sr. Director, Methodology & AI Innovation, Analytics & Technology, IQVIA
TOM HAYES, Practice Leader, Brand Analytics & Technology, IQVIA

Across life sciences, every company has been hard at work trying to put AI at the center of their analytics strategy. After many months of strategic initiatives, investments, and pilot projects, successfully scaling AI across the enterprise is proving to be an enormous uphill climb. The journey began with traditional analytics, moved to integrated advanced analytics, and now the target is an insights-driven enterprise that delivers at least 10-15% revenue uplift and 20-30% cost reduction compared to previous methods. Organizations have invested heavily in analytics platforms, advanced models, and pilot use cases designed to enhance productivity and performance. Yet, despite this momentum, successfully scaling AI is still out of reach for many large and small companies. Most initiatives stall in experimentation, delivering pockets of value without fundamentally changing how the organization operates.

This challenge is not the result of insufficient ambition or immature technology. It's a consequence of how AI has historically been implemented: as an extension of analytics rather than as an operating model. To unlock real enterprise impact, pharma must move beyond insights generation toward AI-enabled execution — where intelligence is trusted and continuously translated into action, at scale.

"Good" in commercial AI and analytics is no longer defined by better dashboards or more sophisticated models. It's defined by an organization's ability to turn **data** into **intelligence** and **actions** from a seamless, governed system that operates at the speed of the business.

5 key shifts shaping the future of AI and analytics

- Data is treated as a high-quality, reusable product with built-in ownership and SLAs
- Transitioning from static dashboards to real-time, conversational analytics
- Moving beyond simple "assistants" to autonomous AI workflow agents
- Shifting skills to AI fluency and human-centric judgement as core leadership assets
- Progress from manual to AI-enforced, real-time governance and compliance

The new requirement: AI-ready data

Fragmented data sources, inconsistent standards, and legacy systems create barriers that even the most advanced algorithms cannot overcome. When data is not discoverable, interoperable, and anchored around real business needs, AI initiatives remain stuck in pilot mode. Preparing data for AI is not a technical cleanup exercise, it's a strategic imperative. Data must be treated as a high-quality, reusable product — with clear ownership, defined service levels, and built-in governance. As AI becomes more autonomous, governance must evolve alongside it. Traditional, manual governance models cannot keep pace with real-time decisioning and execution. Notably, 36% of organizations say their data is either "mostly insufficient" or only "moderately sufficient" to support AI, highlighting a persistent gap in data readiness. Just over half (52%) rate their data as "mostly sufficient," and only 11% say it is "fully sufficient".[1]

1. https://www.iqvia.com/library/white-papers/ai-in-life-sciences-commercialization

Enterprise AI requires governance that is embedded, automated, and enforced in real time. Policies, constraints, and compliance rules must be built directly into data pipelines, models, and agents, ensuring that AI operates within defined boundaries without slowing the business down. This shift represents a major change in mindset. Governance is no longer something that happens after the fact. It becomes part of how intelligence flows through an organization. Without this, organizations spend the majority of their time reconciling data rather than acting on it. Enterprise AI depends on data that is not just clean, but contextualized — data that reflects how the business actually operates, how decisions are made, and how outcomes are measured and optimized.

As pharma organizations scale analytics and interconnected autonomous AI capabilities, a foundation of healthcare-grade, AI-ready data is non-negotiable. Domain-trained models are required to achieve the level of trust and accuracy necessary in life sciences. AI-ready data is information that is explicitly structured, validated, and contextualized to enable AI systems that do more than analyze — they act. Autonomous and semi-autonomous agents cannot rely on loosely reconciled spreadsheets or inconsistent hierarchies. This shift raises the bar. Automation does not tolerate ambiguity. If inconsistencies exist in customer records, product definitions, or organizational rules, AI will surface the inconsistencies, and even amplify them.

AI-ready data is required to engineer trusted intelligence at scale and to consistently meet privacy and security standards for life sciences. It allows AI to move from recommendation to execution without introducing risk. Much of today's AI in pharma still takes the form of "assistants" — tools that help users analyze data, generate content, or answer questions. Field engagement agents are reducing call prep from hours to minutes. While valuable, assistants alone do not fundamentally change how work gets do. The next phase of enterprise AI is agent-based with digital agents designed to unify complex data streams, understand domain-specific rules, and coordinate action across systems.

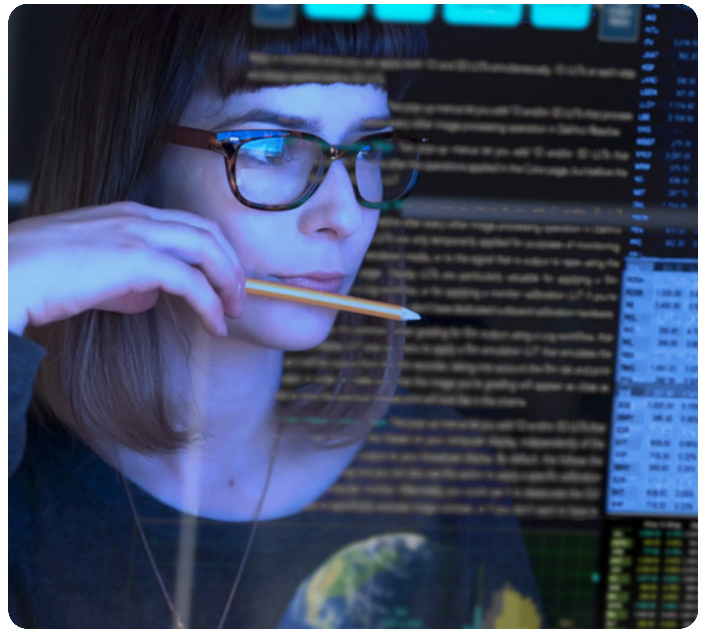

Intelligence at scale

Sales, marketing, analytics, data science, and operations often operate within parallel ecosystems, each with its own data models, dashboards, tools, and vendors. These environments may be individually sophisticated, but collectively they create invisible silos that slow the organization down. At the center of AI-ready data sits Master Data Management (MDM). Once viewed as a back-office discipline, MDM has become a critical enabler of enterprise AI. MDM establishes a single, trusted view of the most important business entities — customers, products, organizations, and relationships — across the enterprise. This consistency is what allows AI systems to operate confidently across functions, channels, and workflows. With a strong MDM foundation, organizations can reduce manual intervention, accelerate data integration, and support complex use cases such as omnichannel engagement, territory design, and regulatory reporting. Changes can be validated once and propagated everywhere, eliminating reconciliation cycles that slow execution.

Commercial teams may rely on one set of reports, while analytics teams maintain separate logic and methodologies. Data scientists develop models that are difficult to operationalize. Field teams receive guidance that is already outdated by the time it reaches them. This fragmentation introduces friction at exactly the moment speed matters most. AI does not struggle because it cannot generate insight; it struggles because insight cannot move efficiently through the organization. What enterprise AI requires instead is a shared foundation — common data, consistent logic, and interoperable systems — that allows intelligence to flow seamlessly from analysis to execution. It doesn't end there; measuring, optimizing performance, and feeding this continuous learning back into insight generation is how enterprise intelligence transforms the way teams work.

Insights to inform strategy and actions

Most commercial teams began their AI journey with traditional analytics — descriptive reporting, historical performance tracking, and standardized dashboards. The evolution of commercial analytics has followed a clear trajectory: from descriptive analytics that explained what happened, to predictive analytics that forecast what is likely to happen, and then to prescriptive analytics that recommend what actions to take. Today, next-generation modeling approaches, such as sequential deep learning, are accelerating this progression by improving model performance and revealing characteristic patterns in patient journeys — dose accelerations, steroid spikes, and shifts in disease activity — that can inform earlier, more targeted interventions. Large Language Models (LLMs) and domain-specific approaches add a reasoning layer that connects these signals with the wider clinical record, including complex, unstructured data, to dynamically support decision-making. As a result, organizations are now moving beyond static recommendations toward intelligent, continuously learning systems that can operationalize decisions in real time.

This foundation then evolved into integrated advanced analytics, where predictive models and more complex analyses supported planning and forecasting. Today, the aspiration has shifted again. AI is enabling insight-driven actions with measurable impact — improving commercial effectiveness, accelerating decisions, driving efficiency, and increasing precision across functions. With 30–40% improvement in forecast accuracy compared to manual or analogue-based approaches, the transformation is real. Advanced methodologies that segment data and model treatment pathways empower teams with richer, faster answers by pinpointing drivers of success with greater granularity and precision.

At the application level, agents do not just inform decisions; they help execute them by generating executive briefings, running analyses, and translating insights into immediate operational changes. Resource allocation adjusts automatically. Campaigns are optimized in near-real time. Field guidance updates dynamically as conditions shift. This is where AI begins to deliver sustained business impact — not by replacing human judgment, but by removing friction from execution.

The challenge goes beyond solving for insight quality and improved workflows. In many organizations, analytics teams already generate highly sophisticated outputs. The real issue is latency and fragmentation: insights are discovered in one system, debated in another, and executed — often manually — in a third, with conditions changing again by the time the action occurs. As long as analytics remain separate from execution, AI's value will remain constrained.

A leading practice that companies should be adopting reflects this shift. Organizations may begin with strong predictive models to identify high-value HCPs and forecast therapy adoption, generating valuable insights but struggle to consistently translate them into action.

Introducing a prescriptive layer such as next best action recommendations around channel, timing, and messaging, is an important step, but on its own is often not enough to drive sustained impact. The next stage is to embed these capabilities within an agent-enabled model, where real-time data signals such as patient starts, formulary changes, and engagement patterns trigger automated, context-aware actions directly within commercial workflows. These systems continuously learn from outcomes, creating a closed-loop model that adapts and improves over time. Companies that move in this direction can drive stronger engagement, respond faster to market dynamics, and unlock measurable improvements in therapy adoption. This demonstrates that the real value of AI lies not just in recommending actions, but in executing and optimizing them at scale.

AI in commercial life sciences becomes transformational when intelligence is embedded directly into workflows, models learn over time, and agile scenario planning enables decisions and actions to happen in near-real time. This is also where AI-driven insights begin to reveal areas for deeper analysis and create the ability to model every patient, HCP, and treatment choice with greater speed and precision.

What "good" looks like

- **Data** is comprehensive, connected, and treated as a reusable enterprise product
- **Governance** is continuous, automated, and built into the operating fabric
- **Intelligence** is domain-trained and embedded into workflows, not isolated in dashboards
- **Action** happens in real time, coordinated across teams, channels, and systems
- **Measurement** is integrated, enabling continuous learning and optimization

In this model, AI is no longer a collection of tools. It is an operating capability and the foundation required to turn intelligence into sustained advantage. Enterprise AI succeeds when organizations stop treating analytics, data, governance, and execution as separate initiatives and start designing them as one system. The move from analytics to autonomy is not incremental. It requires rethinking how data is managed, decisions are executed and optimized, how models continuously learn, and how humans and machines work together. For organizations that build these foundations now, the payoff is significant: faster decisions, more responsive execution, and measurable impact at scale. For those that do not, AI will remain perpetually promising — and perpetually out of reach.

IQVIA enterprise AI and analytics

Connected commercial data

- Enterprise-wide comprehensive and connected view of all life sciences entities and deep insights about them
- Fragmented information optimized into AI-ready data

Intelligence

- Domain-trained AI/ML models turning data into intelligence interface that delivers actionable insights at scale
- Transforming how logic engines are developed to inform execution

Action

- Real-time coordination of insights into executable actions across commercial teams and engagement channels
- Interoperable agents that connect insights to an actionable workflow
- Embedded continuous measurement to inform learning and enhance performance outcomes

IQVIA

CONTACT US

iqvia.com