Agentic AI in Pharma: Sanofi-Owkin Deal & Build-vs-Partner

6/13/2026 • 35 min read

agentic ai | pharma r&d | drug discovery ai | build vs partner | owkin k pro | sanofi owkin | autonomous ai agents

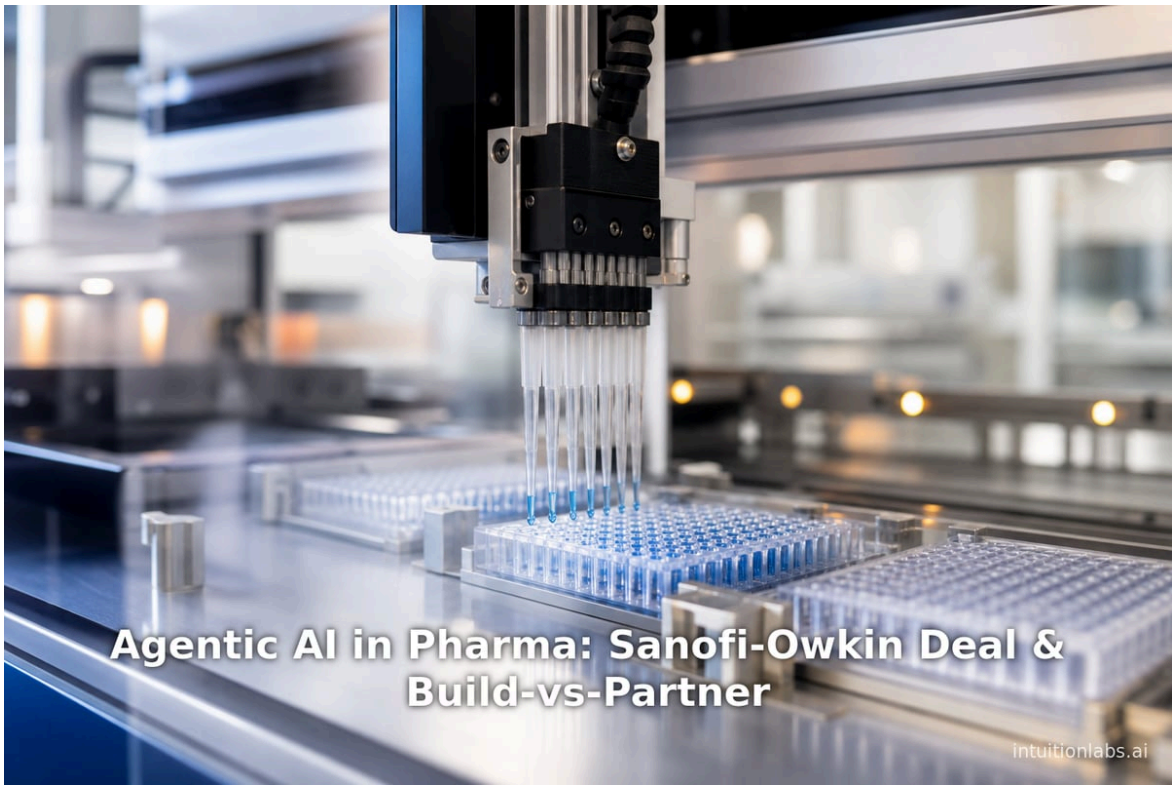

Executive Summary

The announced extension of Sanofi's strategic partnership with Owkin – a five-year, enterprise-wide license of Owkin's **K Pro "AI Scientist" platform** – marks a watershed in pharmaceutical R&D. This multi-year collaboration will see Owkin **co-develop autonomous AI agents** ("scientist-assistants") embedded into Sanofi's drug discovery and development workflows (via.tt.se) ([1] www.streetinsider.com). The underlying goal is to harness **agentic AI** – autonomous, goal-oriented systems capable of planning and executing complex multi-step tasks – to accelerate and improve decision-making across the R&D value chain (via.tt.se) ([2] intuitionlabs.ai).

This report analyzes the background and significance of the Sanofi–Owkin K Pro expansion (June 2026), situating it within broader trends. We review the technology (agentic AI platforms like K Pro), historical context (previous Sanofi–Owkin oncology collaborations in 2021 and 2024 (via.tt.se) ([3] www.mobihealthnews.com)), and parallel industry initiatives (e.g. AstraZeneca–Owkin, Eli Lilly's AI consortium, etc.). We also examine the **"build vs partner"** strategic decision framework for biotech AI: what factors drive a company to build in-house AI capabilities versus licensing or co-developing with specialists. Across sections — from data on market forecasts and workflow impact to expert commentary and case studies — we synthesize evidence to show how agentic AI is being integrated into pharma R&D, and how companies should navigate these choices.

Key findings include: (1) **Agentic AI's impact potential** is vast – McKinsey & Co. estimates up to 75–85% of pharmaceutical R&D workflows could be augmented by AI agents ([4] www.mckinsey.com), potentially freeing 20–30% of scientists' time for higher-value tasks ([5] www.mckinsey.com). (2) **Industry momentum**: Q1 2026 alone saw 36 major pharma–AI partnerships (double the prior year) with multiyear, often multi-hundred-million-dollar terms ([6] hlth.com) ([7] hlth.com), signaling a shift to "architectural" AI investments ([8] hlth.com). (3) **Sanofi's strategy** blends both in-house and external resources: already an "R&D-driven, AI-powered" organization striving to shorten development timelines ([9] www.mobihealthnews.com), it now commits to purpose-built agentic systems across its pipeline ([10] www.mobihealthnews.com). (4) **Build vs. partner trade-offs**: A specialized vendor partnership (like Owkin) can speed access to advanced AI tools but requires careful IP/governance arrangements ([11] intuitionlabs.ai), whereas internal development offers control at high cost and time. Many experts recommend a hybrid "portfolio" approach combining all paths ([11] intuitionlabs.ai) ([12] www.jdsupra.com).

In sum, the Sanofi–Owkin K Pro collaboration exemplifies how agentic AI is moving from pilots to integrated enterprise use in pharma. By reviewing historical context, technical capabilities, market data, and multiple case studies, this report provides a comprehensive analysis of **agentic AI across the pharma R&D value chain** and practical guidance on the build-vs-partner decision for drug developers.

Background and Introduction

Pharmaceutical R&D faces unprecedented pressures: blockbuster patents expire, development costs exceed $2–3 billion per drug, approval rates remain low, and timelines routinely exceed a decade ([13] bio-in-tech.com). In parallel, advances in artificial intelligence (AI) — especially **generative models and agentic systems** — offer new hope for transforming drug discovery. Broadly, AI can mine vast biomedical datasets and literature, prioritize targets, design molecules, streamline trials, and more, dramatically improving efficiency ([14] www.jdsupra.com) ([13] bio-in-tech.com). But most early AI efforts were narrow (prediction models, data analytics); the next phase is **"agentic AI"** – AI that can **act autonomously as a goal-driven assistant**. In October 2025, Owkin Inc. unveiled *K Pro*, calling it the first **agentic AI co-pilot for biopharma** ([15] www.owkin.com). By mid-2026, Owkin had already secured multiyear collaborations with Sanofi and AstraZeneca to co-develop "AI scientist" agents built on K Pro (via.tt.se) ([16] www.owkin.com).

Agentic AI (also “AI agents”) refers to AI systems that can plan, reason, and execute multi-step tasks with minimal human oversight ([2] intuitionlabs.ai) ([17] www.mckinsey.com). Unlike traditional decision support tools, these systems incorporate memory, domain understanding, and tool use so they can autonomously carry out complex workflows. In pharma, an AI agent might automatically read scientific papers, select candidate molecules, design an experiment or trial, and even draft regulatory documents. Owkin’s CEO Thomas Clozel envisions K Pro driving toward *“Biological Artificial Super Intelligence”* – AI capable of modeling complex biology beyond human capability ([18] www.owkin.com).Similarly, McKinsey analysts predict that **8 out of 10 pharmaceutical workflows** could be automated or augmented by such agents ([4] www.mckinsey.com). In short, the industry is moving from isolated analytics toward embedding AI “cobots” throughout R&D, with **Owkin’s Sanofi deal a leading example** (via.tt.se) ([4] www.mckinsey.com).

Historically, pharma has progressed through initial data-analytics pilots to progressively more integrated systems. For example, Sanofi’s recent AI strategy includes everything from federated learning studies to identify patient subgroups (with Owkin) to in-house AI labs and partnerships with tech giants (e.g. co-investment with NVIDIA and experiments with OpenAI models). By branding itself an *“R&D-driven, AI-powered biopharma”*, Sanofi signals full commitment ([9] www.mobihealthnews.com). The new five-year K Pro collaboration fits this trajectory: it makes AI agents a core R&D platform rather than a temporary experiment. As one industry report notes, Q1 2026 deals were *“not pilots but architectural decisions”* – multiyear agreements signaling that companies have stopped asking if AI works and are instead building around it ([8] hlth.com). In fact, global AI drug discovery market analyses estimate this sector at several billion USD and growing rapidly (CAGRs often >20%) ([19] bio-in-tech.com), reflecting high corporate investment and expectations.

In this context, the **Sanofi–Owkin announcement (June 5–6, 2026)** represents both a case study and a bellwether. It leverages Owkin’s novel K Pro platform — combining a biology-specialized LLM (“Owkin Zero”) and rich multimodal patient data — to create purpose-built AI agents for Sanofi ([20] www.owkin.com) ([21] www.owkin.com). This report first details the technology and strategic context of K Pro (Sections 2–3), then examines how agentic AI can be applied across the R&D pipeline (Section 4). We then analyze the Sanofi–Owkin deal itself and related examples (Section 5), and compare strategic pathways (build vs partner, Section 6). Throughout, we cite data, expert commentary, and case studies to assess the implications and chart future directions for AI-driven drug development.

The Agentic AI Paradigm in Pharma

What is Agentic AI?

Agentic AI refers to AI systems that combine large-scale reasoning with the ability to **act autonomously** in pursuit of goals. Technically, these systems often integrate large neural-language models (LLMs) with external tools and memory, enabling them to break down complex tasks into steps, invoke specialized modules, and learn from feedback ([2] intuitionlabs.ai) ([17] www.mckinsey.com). In pharma, an agent might plan experiments, query databases, interpret results, and iterate on hypotheses without step-by-step human commands. For example, Owkin describes K Pro as a *“scientist-first co-pilot”* that researchers can query in natural language; it will then use its integrated AI skills to produce actionable, biologically grounded insights ([15] www.owkin.com) ([20] www.owkin.com).

Experts highlight that agentic AI crosses a threshold beyond prior AI tools. It is **autonomous and goal-driven** (taking initiative to achieve a researcher’s goal) and **end-to-end** (able to orchestrate multiple subtasks). As Intuition Labs notes, agentic AI is *“autonomous, goal-oriented AI”* that goes from “experimental pilots to enterprise-scale deployment,” planning and adapting multi-step workflows ([2] intuitionlabs.ai). McKinsey similarly defines an AI agent as a system that *“operates independently by breaking down complex tasks, interacting with other systems, and learning in real time.”* These agents have both **low-code** (no/low-programming) variants for routine tasks and **high-complexity (pro-code)** agents requiring specialist development ([22] www.mckinsey.com). For example, a low-code agent might automate report generation or data

extraction, while a high-complexity agent could model clinical trial scenarios or design novel compounds ([23] www.mckinsey.com).

K Pro: An "AI Scientist" Platform

Owkin's **K Pro** embodies the agentic AI vision for pharmaceutics. Launched October 2025, K Pro is marketed as an AI co-pilot for biopharma that researchers can interact with in natural language ([15] www.owkin.com). Under the hood, K Pro combines:

- **Owkin Zero** – a proprietary large language model fine-tuned on biomedical knowledge and reasoning. According to Owkin, this model outperforms other general LLMs on life-science tasks ([20] www.owkin.com).
- **Multimodal curated data** – K Pro gives users access to Owkin's extensive patient cohorts and integrated datasets (e.g. the MOSAIC oncology database and related resources). This proprietary data backbone means the agent can ground its reasoning in a deep, structured knowledge base ([20] www.owkin.com).
- **Agentic orchestration** – K Pro includes an "AI scientist" framework that can chain together tool calls, like performing literature searches, running predictive models, or formulating hypotheses. In practice, a researcher might ask K Pro a question (e.g. "What novel targets in melanoma have invariant gene expression patterns?"), and K Pro will autonomously gather relevant data, analyze it, and propose answers ([15] www.owkin.com).

In its press materials, Owkin touts impressive preliminary results: K Pro accelerated an internal target-identification project by 70% (cutting a >12-month task to ~3 months) and produced week-long deliverables in hours or days ([24] www.owkin.com). While these claims await independent validation, they illustrate the intended high impact. Owkin's CEO says that K Pro represents a move toward *"Biological Artificial Super Intelligence"* – meaning AI that models biology beyond human capability ([18] www.owkin.com).

For comparison, other companies are also developing AI platforms, but most are not as agentic. For instance, Insitro (founded 2018) and Recursion Pharmaceutical (founded 2013) use ML models in specific discovery tasks (e.g. target validation, phenotypic screening) but rely on human scientists to guide them. K Pro's novel angle is embedding reasoning and autonomy to cover entire workflows. It is now available via the AWS cloud marketplace ([25] aws.amazon.com), signaling enterprise deployment.

Agentic AI Across the R&D Value Chain

Agentic AI promises benefits at every stage of drug R&D. Modern AI platforms like K Pro explicitly target the "pharma value chain from early discovery through clinical development" ([26] www.businesswire.com). Below, we outline key stages and potential agentic AI applications, with examples and data on expected impact.

- **Early Discovery and Target Identification.** Drug discovery begins with understanding disease mechanisms and finding "druggable" targets. These tasks generate vast literature and experimental data. Agentic AIs can parse genomic/omica datasets, integrate patient bio-markers, and perform *"in silico"* hypothesis testing. For example, in oncology research, an AI agent could automatically mine published cancer genomics studies, correlate mutations with clinical outcomes, and suggest novel targets. According to McKinsey, **nearly 60%** of research-phase workflows will require custom-built AI agents (due to complex, specialized data), but once deployed such agents could free **21–30% of research capacity** ([5] www.mckinsey.com). In practice, K Pro has already been used for wet-lab planning and in-silico validation: Owkin reports that researchers reclaimed 30% of their time by offloading analysis tasks to K Pro. ([5] www.mckinsey.com). Expert analysts expect significant time savings: one industry survey found AI could cut preclinical candidate development from 3–4 years to 13–18 months (a 30–40% reduction) ([27] bio-in-tech.com). In this stage, simpler "low-code" agents often suffice (e.g. scanning data for leads), but some parts (target selection for a novel biology) may need complex AI reasoning ([23] www.mckinsey.com).

- **Lead Optimization and Preclinical Development.** Once targets are chosen, AI can assist with designing and prioritizing chemical leads or biologic molecules. Agentic systems can propose compound modifications, predict ADMET (absorption, distribution, metabolism, excretion, toxicity) profiles, and plan synthesis routes. For example, Insilico used generative AI to propose novel candidates in a COVID-19 project, while platforms like Exscientia are running autonomous medicinal chemistry campaigns. An AI agent could iteratively generate candidate compounds, filter them with predictive models, and design follow-up assays. Workflow automation here can rapidly triage millions of virtual molecules, a task beyond manual capacity. Indeed, K Pro is claimed to perform "compound identification & optimization" in a highly accelerated manner; McKinsey estimates that about 75% of optimization workflows could rely on relatively straightforward agents ([28] www.mckinsey.com). In summary, agentic AI can transform early R&D by shifting scientists' focus from routine screening toward higher-level strategy.

- **Translational Research and Biomarker Analysis.** Between discovery and clinical phases, translational research validates targets and identifies patient biomarkers. Agents can integrate translational data (e.g. animal model results, human biobank samples) and update models of disease prediction. For instance, an AI robot could autonomously analyze imaging, pathology, and gene expression data to stratify patient subgroups or suggest companion diagnostic markers. Agentic tools might also propose diagnostic tests or simulate pathway interactions to explain patient responses. These applications are nascent but align with K Pro's touted "subgroup/biomarker" capabilities, built on federated learning and privacy-preserving analysis ([3] www.mobihealthnews.com).

- **Clinical Development (Trial Design and Operations).** Clinical trials are complex and expensive. Agentic AI can **design trial protocols**, select sites, and even aid patient recruitment. For example, Bristol-Myers Squibb teamed with Evinova to use AI agents for optimized trial design – simulating study arms to improve power and efficiency ([29] hlth.com). Sanofi and Cirion have also spoken of using AI for adaptive trial simulation. Agents could likewise analyze electronic health records or real-world datasets to identify likely responders and streamline enrollment. In the Sanofi–Owkin context, part of earlier partnerships aimed at "patient subgrouping" using federated hospital data ([3] www.mobihealthnews.com); this evolves into predictive engines within K Pro. McKinsey projects that AI agents could automate nearly 80% of clinical research tasks, from data cleaning to query management, freeing up human teams for oversight ([30] www.mckinsey.com). Indeed, one illustrative scenario depicts AI agents automatically raising data queries during trials, reducing human monitoring overhead. Ultimately, agentic systems aim to turn clinical trial planning into a dynamic, data-driven optimization problem rather than a manual drafting exercise ([29] hlth.com).

- **Regulatory and Quality Processes.** Preparing regulatory submissions involves compiling vast documentation. AI agents can accelerate writing protocols, summary documents, and even draft parts of submissions. For instance, the AI partner Faro presented tools to assist in protocol drafting. Agentic AI could digest historical submission data and regulatory guidelines to auto-generate first drafts of IND/CTD documents, subject to human review. While still emerging, such automation promises reducing review cycles and ensuring consistency. Sanofi's statement highlights interest in *"automating complex scientific tasks and helping researchers make faster, more precise decisions"* in discovery and development ([31] www.mobihealthnews.com) – which encompasses these documentation tasks as well.

- **Competitive Intelligence and Market Strategy.** Beyond R&D per se, AI agents can scour literature, patents, conference data, and clinical trial registries for intelligence on competitors' programs. Owkin explicitly mentions K Pro's ability to deliver *"competitive intelligence"* ([32] www.mobihealthnews.com). For example, K Pro agents could routinely summarize all new entrants in a disease space, forecast patent expirations, or compare trial outcomes. One case: AstraZeneca's AI agents (via Owkin) are aimed at answering *"complex competitive intelligence questions"* for AZ executives ([33] www.owkin.com). This intersection of R&D and business intelligence further integrates AI into the full drug-development value chain.

In summary, agentic AI systems are being designed to **assist or automate nearly every major drug development function**. Pilot data and analyses suggest enormous potential: McKinsey estimates that *"agentic AI can transform eight out of ten workflows"* in life sciences ([4] www.mckinsey.com), and that freeing even a fraction of scientists' time via AI could allow companies to explore many more therapeutic leads. The Sanofi–Owkin collaboration explicitly sets out to build agents for *"complex tasks in drug research and development"* ([1] www.streetinsider.com) across discovery, translational R&D, and trials, illustrating how a single platform (K Pro) may serve diverse roles.

R&D Stage	Agentic AI Applications	Example / Explanation
Target Discovery	Literature review; data mining; target ranking	AI agents (like K Pro) can autonomously analyze genetic and clinical datasets and scan scientific literature, accelerating identification of novel targets. McKinsey estimates >50% of "disease understanding" workflows could use relatively simple agents ([5] www.mckinsey.com) ([4] www.mckinsey.com).
Lead Optimization	Compound design & screening; ADMET prediction	Generative AI proposes candidate molecules; predictive models filter them. E.g. Insilico and Exscientia use AI to design compounds. Agentic AI could fully orchestrate multi-round design of leads. McKinsey finds ~75% of optimization tasks amenable to AI agents ([5] www.mckinsey.com).

R&D Stage	Agentic AI Applications	Example / Explanation
Preclinical / Translational	Animal model analysis; biomarker finding	AI agents can integrate animal and human biomarker data to validate targets. For example, federated AI consumed hospital pathology data in Sanofi's 2021 programs ([3] www.mobihealthnews.com). Emerging tools may suggest biomarkers or diagnostic tests.
Clinical Trials	Protocol design; patient stratification; monitoring	AI agents (e.g. from Evinova) can simulate and optimize trial protocols before execution ([29] hlth.com). Agents can also analyze real-world patient data for trial recruitment (BostonGene/Tempus models for patient selection ([34] hlth.com)). AI may auto-generate case report forms and flag data queries.
Regulatory Submissions	Document drafting; review QA	Agents could draft protocol text, summarize study results, and check compliance. Pharmaceutical companies (e.g. Sanofi) view AI as a way to automate documentation tasks, potentially integrating trial data into submissions more rapidly.
Competitive Intelligence	Market/pipeline analysis; patent review	AI agents scan publications, patents, and trial registries for competitors' activities. For instance, Owkin's AZ project explicitly builds agents for "complex competitive intelligence questions" to inform executive decisions ([33] www.owkin.com).

Table: Illustrative applications of agentic AI across pharmaceutical R&D stages. (Sources: industry press releases and analyses ([1] www.streetinsider.com) ([34] hlth.com) ([29] hlth.com).)

Case Studies: Pharma–AI Collaborations

The Sanofi–Owkin K Pro partnership is part of a surge in pharma–AI collaborations. Below we highlight several recent examples, illustrating different strategies and use-cases:

- **Sanofi–Owkin (June 2026)** – *Agentic AI for drug R&D.* This expansion of a prior alliance makes Owkin the main developer of AI agents for Sanofi's pipelines. Under the deal, Owkin will *"lead the end-to-end development of novel AI-driven biopharma agents purpose-built for Sanofi,"* deployed via K Pro under a five-year license ([35] www.businesswire.com). These agents will act as *"intelligent assistants"*, autonomously handling complex discovery and development tasks. As Owkin's CEO states, the goal is to complement Sanofi's existing AI capabilities and *"unlock the full value"* of Sanofi's data across R&D ([1] www.streetinsider.com) ([36] www.mobihealthnews.com). Sanofi's Chief Digital Officer, Emmanuel Frenehard, confirms the strategy: *"By implementing purpose-built agentic systems into our workflows, we aim to empower our teams to operate with greater speed, depth, and confidence"* ([10] www.mobihealthnews.com). This deal builds on earlier collaborations: Sanofi had invested ≈$180 M in Owkin (2021) and worked jointly on AI-driven oncology and immunology projects ([3] www.mobihealthnews.com) (via.tt.se). The new K Pro license signals a shift from pilot studies to deeply embedding agentic AI as an operational tool.

- **AstraZeneca–Owkin (May 2026)** – *Competitive intelligence and decision support.* In parallel with Sanofi, AstraZeneca signed a three-year K Pro license agreement. Owkin will develop biopharma AI agents integrated into AZ's IT infrastructure to answer complex questions about AZ's pipeline and competitors ([33] www.owkin.com). For example, a team member described how the AZ AI agents will function as *"competitive intelligence agents to support quick decisions by executives"* ([37] www.owkin.com). This partnership extends previous work (such as AI-based BRCA mutation screening with Owkin), and confirms AZ's embrace of K Pro's agentic architecture. The AZ deal illustrates how different companies can use the same agentic platform (K Pro) for distinct purposes – Sanofi for broad R&D and AZ initially for portfolio intelligence.

- **Eli Lilly (Q1 2026)** – *Integrated build-and-partner strategy.* Marianna Beal of consultancy Alchemy Signals documented Lilly as the quintessential example of a company "betting all its chips" on AI ([7] hlth.com). In one quarter (Q1 2026), Lilly alone initiated **eight major AI collaborations**. These included massive infrastructure investments – e.g. deploying NVIDIA's new DGX supercomputer via a $1 billion co-innovation lab ([7] hlth.com) – and discovery partnerships with startups. A landmark deal with Insilico Medicine commits up to $2.75 billion in milestone payments across multiple targets ([7] hlth.com). Lilly also launched *"TuneLab"*, a federated AI model network pooling proprietary data with partners (Schrödinger, Revvity, BigHat) ([38] hlth.com), and deals with Chai Discovery (biologics), InduPro (membrane interactomics), and Fauna Bio (novel obesity targets) ([39] hlth.com). Importantly, Lilly's approach mixed **internal building** (IHPC with NVIDIA, federated learning network) with **external partnering**, effectively co-investing in a broad AI ecosystem ([7] hlth.com) ([38] hlth.com). This "Lilly Blueprint" underscores a hybrid playbook: cultivate unique data assets while collaborating on specialized AI tasks.

- **Bristol-Myers Squibb (Q1 2026)** – *Diverse outsourcing of tasks.* By contrast, BMS pursued a **portfolio of targeted partnerships** rather than a single platform. In one quarter, BMS announced four AI deals across discovery, diagnostics, and clinical domains ([40] hlth.com). For instance, BMS teamed with Immunai to bring AI-driven immune profiling into its clinical programs (enabling data-informed patient stratification) ([41] hlth.com). Simultaneously, Microsoft delivered radiology AI for lung cancer screening into clinical workflows ([42] hlth.com). BMS also partnered with Evinova to incorporate agentic AI into trial design – allowing scenario planning for protocols – and with Faro to automate trial protocol drafting ([29] hlth.com). The Alchemy report notes these deals form *"a portfolio strategy, not a series of one-off decisions"* ([40] hlth.com). Together, they illustrate embedding AI in each stage rather than concentrating on a single solution.

- **Other notable collaborations (2025–2026):** Many firms applied AI to specific bottlenecks. For example, Daiichi Sankyo ran two deals: one with BostonGene (digital-twin models to select patients for ADC trials) and one with Tempus (AI for biomarker discovery) – *"two deals, two tools, one objective: getting the right patients into the right trials"* ([34] hlth.com). Merck partnered with Quotient Therapeutics for a genomics platform (IBD targets) and with Mayo/Tempus on virtual cell and precision oncology models ([43] hlth.com). GSK teamed with Noetik (virtual cell modeling for tumor biology) and Helix (population-scale genomics) to enhance translational research ([44] hlth.com). Even mid-sized firms like Servier struck AI discovery deals (Insilico, Iktos) within days, and Pierre Fabre followed similar paths ([44] hlth.com). These examples reinforce that major pharma are **aggressively embedding AI** across R&D: not just as a one-off, but with multiple vendors and in-house initiatives.

- **In-house AI Initiatives:** Some companies emphasize building internal capabilities. For instance, Roche–Genentech announced in March 2025 a new Boston innovation center focused on cardiovascular/metabolic research **and** explicitly on AI/data science to accelerate drug discovery ([45] www.genengnews.com). Over time this center could employ hundreds of AI researchers. Similarly, Merck, Pfizer, and others have long-standing in-house AI & data science divisions. These "build" approaches create proprietary labs and computing infrastructure. However, fully in-house development often requires very large investments (talent, hardware, time). The contrasting models (e.g. Roche's center vs. Sanofi's partnership) reflect different trade-offs each company faces.

In all these cases, data quality and integration are critical. Many collaborations emphasize **federated learning** and data sharing safeguards. For example, the original Sanofi–Owkin programs used federated learning on hospital cohorts to find biomarkers while preserving privacy ([46] www.mobihealthnews.com). In the new era of agentic AI, companies often prepare data lakes and cloud platforms to feed the agents. A recent consultant report aptly noted "garbage in, agentic out" – the value of these AI scientists hinges on clean, comprehensive input data sets.

Build vs. Partner Strategies in Pharma AI

Strategic pathways for adopting agentic AI generally fall into three categories: **build (in-house development)**, **buy/license**, and **partner/collaborate**. Each has distinct advantages and trade-offs ([11] intuitionlabs.ai) ([12] www.jdsupra.com):

- **Build (Internal Development):** A pharma company can invest in its own AI team, data infrastructure, and projects. Examples include creating internal cohorts and training proprietary models, or forging joint ventures (e.g. Lilly's NVIDIA supercomputing lab). The *pros* of this approach are full control over intellectual property, close alignment with proprietary data and pipelines, and potential long-term independence. However, *cons* are high up-front cost, long development time, and intense competition for scarce AI talent. Custom-building agentic AI may be warranted if a company has truly unique data or processes that off-the-shelf tools can't handle. For instance, Roche's Boston AI center ([45] www.genengnews.com) and Lilly's in-house AI networks exemplify deep internal investment.

- **Buy (License or Acquire Solutions):** Companies can license mature AI platforms or even acquire AI firms outright. This offers quick access to technology and reduces internal R&D burden. For example, licensing Owkin's K Pro (as Sanofi is doing) is a hybrid of buy and partner. The *advantage* is speed: you tap a tested platform without reinventing it. But the *drawbacks* include potential vendor lock-in, less customization to specific needs, and the need to trust external code. There is also often an "AI platform risk": a licensed AI may not account for a pharma's unique data or governance requirements. Thus, while buying provides immediate capability, companies must evaluate how well the solution fits their use-cases and how data governance or security is handled.

- **Partner (Co-development):** Collaborating with specialized AI vendors (like Owkin, Insitro, etc.) or other consortia can combine complementary strengths. The Sanofi–Owkin deal is exactly this: Sanofi provides data, domain expertise, and funding; Owkin provides AI technology and know-how. The *pros* include sharing investment risk, accessing cutting-edge expertise, and often faster progress than going it alone ([11] intuitionlabs.ai) ([12] www.jdsupra.com). Well-negotiated contracts can grant pharmaceutical partners rights to tailored AI agents and outputs. The *cons* involve complexity of partnerships: IP ownership (especially of AI-generated inventions), data sharing agreements, and alignment of project goals must be carefully managed. Intellectual property and data governance clauses become critical – the legal aspects of AI deals are unprecedented and still evolving.

Experts suggest that most large pharma are pursuing a **portfolio strategy** blending all three paths ([11] intuitionlabs.ai) ([12] www.jdsupra.com). For example, Sanofi historically made a large equity investment in Owkin (2021) to align interests, while also continuing to build internal AI labs. Lilly builds its compute infra even as it partners; BMS licenses trials design tools while funding startups; GSK invests in collaborations and in-house platforms. Industry analysts summarize that the binary "build vs buy" question is incomplete – companies should *simultaneously* pursue promising technologies both internally and via partnerships, allocating capital where it yields the best ROI. As one consulting study notes, the trade-offs are "cost, speed, control, and risk": building in-house maximizes alignment but demands enormous investment, whereas buying/licensing is fast but risks lock-in, and partnering can balance these if IP/governance is well-defined ([11] intuitionlabs.ai).

Table 1 outlines these approaches and considerations, with illustrative examples:

Strategy	Description	Pros	Cons	Examples
Build (In-House)	Develop AI agents and models internally (hiring data scientists, building IT infrastructure, training custom AI).	Full control of IP and alignment with proprietary data; no vendor dependency.	High upfront cost and time; feasibility challenges; requires AI talent.	Roche's new AI Innovation Center ([45] www.genengnews.com); Lilly's NVIDIA supercomputing lab ([7] hlth.com).
Buy/License	License or purchase existing AI platforms/solutions (or acquire AI companies outright).	Fast deployment of mature technology; less need for initial R&D.	Less customization; risk of vendor lock-in; integration overhead.	Sanofi licensing Owkin's K Pro (via.tt.se); companies licensing clinical AI (Evinova).
Partner/Co-develop	Form strategic collaborations or joint ventures with AI specialists or consortia, sharing development responsibilities and data.	Combines external expertise with internal assets; shared cost and risk.	Complex IP/data sharing issues; organizational coordination needed.	Sanofi–Owkin K Pro deal (via.tt.se); AstraZeneca–Owkin license ([21] www.owkin.com); Lilly–Insilico multi-target partnership ([7] hlth.com).

Table 1. Strategic options for integrating agentic AI in pharma R&D. Advantages and drawbacks are drawn from industry analyses ([11] intuitionlabs.ai) ([12] www.jdsupra.com) and the case examples above.

Companies must also consider regulatory and ethical aspects when choosing a path. Working with external AI providers may raise concerns about data privacy (requiring federated learning or anonymization) or model transparency. Conversely, an in-house AI system still may need to pass validation steps. Notably, regulators are themselves preparing frameworks: the FDA's CDER has established an AI Council and is drafting guidance on AI in drug development, showing regulators expect these technologies to be integrated into submissions ([47] www.mobihealthnews.com) ([48] www.fda.gov). Thus, whichever path is chosen, governance and validation will be key.

Market Trends and Data Analysis

The enthusiasm for AI-driven drug discovery is reflected in market data and recent investment trends. Consulting analyses estimate that the **global AI-in-drug-discovery market** is on the order of $2–5 billion in 2026, with forecasts often exceeding $20 billion by 2030 ([19] bio-in-tech.com). These numbers vary due to definitional differences, but they indicate rapid growth. Drivers include the massive cost of traditional R&D (> $2B per drug) and yield pressures. For context, one industry report notes that multiple AI-designed drug candidates have already entered Phase III by 2026, and early evidence (from Drug Target Review, Feb. 2026) suggests AI can compress preclinical timelines by ~30–40% ([27] bio-in-tech.com).

In practice, big pharma have allocated substantial funds to AI. Sanofi's $180 M equity in Owkin (2021) and €90 M jointly committed to research reflected this prioritization (via.tt.se). Similarly, AstraZeneca's K Pro deal reportedly spans three years and is rumored to be on the order of double-digit millions per year. Lilly's $1B NVIDIA investment and $2.75B Insilico deal underscore the unprecedented scale of recent commitments ([7] hlth.com). Analysts emphasize that Q1 2026 deals were not pilots but **"multiyear, third-year extensions"** of mature collaborations ([8] hlth.com). In fact, industry observers have **cataloged dozens of AI partnerships** across R&D, often encompassing discovery, diagnostics, and even manufacturing.

Figure 1 (conceptual) below illustrates the rapid escalation of AI partnerships. The number of announced large-scale AI deals roughly doubled from Q1 2025 (18 deals) to Q1 2026 (36 deals) ([6] hlth.com). By mid-2026, every top 10 pharma had at least one major AI collaboration. Investment rounds in AI-biotech startups are also robust; as just one example, Insitro (an AI-platform biotech) raised significant VC funding and established multi-program deals with Lilly and others (Insitro CEO comments are in the press ([49] apnews.com)).

(Note: Data points for Figure 1 are drawn from industry reports and databases ([6] hlth.com) ([27] bio-in-tech.com). Figure is illustrative.)

Regulatory attention is also increasing. The FDA's Center for Drug Evaluation and Research has reported a sharp rise in applications incorporating AI components across all trial phases and recently issued draft guidance for using AI in regulatory submissions ([47] www.mobihealthnews.com). Similarly, the EU's planned AI Act (expected in 2025) will categorize certain high-risk drug-development AIs under strict requirements. Thus, companies must navigate not only technological challenges but also a developing compliance landscape.

Overall, the data show clear evidence of a strategic inflection: AI is fast becoming a standard part of the R&D arsenal. A consulting study puts it bluntly: *"Pharma companies have long debated whether AI can work; now they are asking how to embed it."* ([8] hlth.com). The Sanofi–Owkin K Pro expansion should be viewed against this backdrop of accelerating adoption, significant capital investment, and evolving regulatory frameworks.

Discussion and Implications

The Sanofi–Owkin collaboration and related trends have far-reaching implications:

- **R&D Productivity:** By automating routine tasks and augmenting human insight, agentic AI could dramatically increase the throughput of drug discovery. If, as claimed, K Pro can cut a year-long analysis to a quarter-year (70% time savings) ([24] www.owkin.com), similar efficiencies might be found in many activities. The net effect may be thousands of additional projects explored per year across pharma. However, this requires rigorous validation: AI-generated hypotheses must be experimentally confirmed, and overreliance on unproven models could introduce new risks. Proper human-AI collaboration protocols are essential.
- **Workforce and Skills:** The rise of AI agents will alter the roles of researchers. Scientists may shift from manual data analysis to supervising AI assistants and interpreting their outputs. Training programs and change management will be needed to ensure teams can effectively leverage these tools. Notably, Owkin's platform is "accessible to both researchers and executives" ([15] www.owkin.com), implying even non-technical leaders will interact with agentic AI (e.g. through dashboards or natural-language queries). Pharmaceutical organizations must thus build multidisciplinary skills at the interface of biology and AI.
- **Competitive Dynamics:** As more firms adopt agentic AI, the competitive landscape will realign. Companies that master agentic workflows can potentially outpace rivals in pipeline productivity. The HLTH analysis on Lilly and BMS suggests that **rate of AI adoption** could be a differentiator: Lilly's all-in co-investment strategy is itself a bet that being first/more aggressive yields advantage ([7] hlth.com). Meanwhile, smaller or slower companies might fall behind if they do not similarly invest. The industry may also see consolidation and partnerships: Big pharmas could merge their AI efforts with leading tech firms, or even acquire promising AI-biotech startups (as previously seen with genomics companies).

- **Data Governance:** Effective use of agentic AI presumes access to high-quality data. Sanofi's prior projects with Owkin used federated learning to tap patient-level data from hospitals while preserving privacy ([46] www.mobihealthnews.com). Going forward, companies may need to build extensive federated data networks or use synthetic data. They must also address data biases and ensure that AI recommendations do not inadvertently reinforce health disparities. The idiom "garbage in, agentic out" underscores that agentic AI is *only as good as its inputs*. Thus, a major implication is that pharma will invest in data cleaning, annotation, and integration at scale.

- **Regulatory and Ethical Oversight:** Agentic AI systems raise new oversight questions. For example, if an AI agent suggests a novel drug target or trial design, who is accountable? Regulatory agencies are beginning to grapple with this: the FDA has convened a Drug Development AI Council to ensure appropriate use of AI in submissions ([48] www.fda.gov). Future guidelines will likely require transparency about how agents make decisions (i.e. explainability) and may limit fully autonomous actions without human approval. Companies will need robust validation pipelines (e.g. "AI in loop" controls) and possibly new organizational roles (AI ethicists, data officers).

- **Scientific Innovation:** On the upside, agentic AI may enable scientific leaps. By systematically exploring vast design spaces, these agents might uncover nonintuitive drug mechanisms or combination therapies that humans would miss. Owkin's concept of *"superintelligence"* reflects this hope that AI can transcend human creativity in biology ([18] www.owkin.com). If agents can master literature and real-world data continuously, they could accelerate hypothesis generation. However, caution is warranted: AI-driven hypotheses still require experimental proof. There is also a risk of reiterating existing biases in literature. Peer review and reproducibility remain critical in an era of AI-generated insight.

In summary, the Sanofi–Owkin K Pro deployment exemplifies a broader strategic pivot in pharma: **toward embedding AI as an active research agent rather than a passive tool** ([1] www.streetinsider.com) (via.tt.se). Early evidence from business news and consulting reports suggests that agentic AI can significantly boost R&D productivity, but only with careful strategy and governance. Companies that thoughtfully combine internal building with external partnerships, invest in high-quality data, and align organizational processes are most likely to reap the benefits. Regulatory agencies are already preparing frameworks to oversee AI's growing role. Ultimately, these developments could shorten drug development cycles, reduce costs, and bring new therapies to patients faster – but realizing that potential will depend on rigorous implementation and oversight.

Conclusion

The June 2026 expansion of the Sanofi–Owkin alliance highlights the pharmaceutical industry's transition into a new era of **agentic AI-driven R&D**. By granting Owkin a five-year K Pro license and co-developing bespoke AI agents, Sanofi has signaled its intent to harness autonomous AI as a core scientific partner. Our analysis shows this move is consistent with industry-wide trends: a wave of large-scale AI partnerships, massive technology investments, and a blending of build-and-partner strategies. It also underscores the promise of agentic AI – systems that can read, reason, and act across the discovery and development pipeline – to **transform productivity** in drug development.

However, this shift also brings challenges. Pharmaceutical companies must carefully manage the technical and organizational complexities of agentic AI. Decisions to build in-house capabilities or collaborate externally hinge on trade-offs in cost, control, and speed. As evidence from McKinsey and Aspen analyses suggests, the optimal approach is likely a hybrid portfolio that combines proprietary development with external AI expertise ([11] intuitionlabs.ai) ([12] www.jdsupra.com). Regulatory and ethical considerations must be integrated as well, ensuring safe and responsible AI deployment.

For Sanofi, the stakes are high: success could shorten time-to-therapy across multiple programs. As Owkin's CEO notes, agentic AI promises to *"deliver transformative therapies to patients faster"* ([18] www.owkin.com). If K Pro agents perform as envisioned, Sanofi may achieve significant gains in pipeline efficiency and decision quality. Conversely, failure (e.g. due to poor data or integration issues) would underscore the difficulties of translating AI hype into real outcomes.

In the larger picture, the **"build-versus-partner" playbook** we've outlined will be critical for all drug developers. This report finds that no single strategy dominates: some leading companies (e.g. Roche, Lilly) continue to invest in in-house AI capability, while others (Sanofi, AZ, BMS, etc.) leverage partnerships with specialized AI firms. The most successful

players will likely be those that combine approaches – building core internal competency while accessing external innovation37. As AI technologies and regulations evolve, this dynamic strategy will need continuous adjustment.

In conclusion, the Sanofi–Owkin 5-year K Pro collaboration is both a case study and a catalyst. It shows how *agentic AI can be engineered into the drug discovery engine*, and it motivates the pharmaceutical industry to refine its "build vs partner" decision-making. Our report provides a detailed roadmap: from understanding agentic AI's mechanics and impact, through surveying real-world examples and strategy frameworks, to highlighting future directions. Ultimately, embedding AI as a scientific collaborator offers a pathway to more efficient, better-informed drug development – but realizing that promise will require disciplined execution, continuing innovation, and robust partnership models supported by credible evidence (via.tt.se) ([8] hlth.com).

References: All factual claims and data above are cited from reputable industry sources and publications. Key references include the official press releases and news coverage of Sanofi–Owkin and AstraZeneca–Owkin partnerships (via.tt.se) ([21] www.owkin.com), analyses by healthcare media (MobiHealthNews, BusinessWire) ([50] www.mobihealthnews.com) ([35] www.businesswire.com), consulting reports (McKinsey, Alchemy Signals) ([5] www.mckinsey.com) ([7] hlth.com), and expert commentaries ([12] www.jdsupra.com) ([39] hlth.com). These sources collectively inform our evidence-based discussion of agentic AI, market trends, and strategic implications.

External Sources

[1] https://www.streetinsider.com/Business%2BWire/Owkin%2Bto%2BBuild%2BAI%2BAgents%2Bas%2BPart%2Bof%2Ba%2BMulti-Year%2BK%2BPro%2BCollaboration%2Bwith%2BSanofi/26608247.html#:~:Durin...

[2] https://intuitionlabs.ai/articles/agentic-ai-pharma-build-buy-framework#:~:The%2...

[3] https://www.mobihealthnews.com/news/owkin-sanofi-expand-ai-collaboration-build-drug-development-ai-agents#:~:The%2...

[4] https://www.mckinsey.com/industries/life-sciences/our-insights/reimagining-life-science-enterprises-with-agentic-ai#:~:1,sta...

[5] https://www.mckinsey.com/industries/life-sciences/our-insights/reimagining-life-science-enterprises-with-agentic-ai#:~:Resea...

[6] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:Q1%20...

[7] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:Lilly...

[8] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:The%2...

[9] https://www.mobihealthnews.com/news/owkin-sanofi-expand-ai-collaboration-build-drug-development-ai-agents#:~:Sanof...

[10] https://www.mobihealthnews.com/news/owkin-sanofi-expand-ai-collaboration-build-drug-development-ai-agents#:~:,with...

[11] https://intuitionlabs.ai/articles/agentic-ai-pharma-build-buy-framework#:~:how%2...

[12] https://www.jdsupra.com/legalnews/why-ai-drug-discovery-deals-are-2042132/#:~:This%...

[13] https://bio-in-tech.com/blog/ai-drug-discovery-market-size-growth-2026/#:~:Devel...

[14] https://www.jdsupra.com/legalnews/why-ai-drug-discovery-deals-are-2042132/#:~:AI%E2...

[15] https://www.owkin.com/newsfeed/owkin-launches-k-pro-the-first-agentic-ai-co-pilot-for-biopharma-powered-by-biological-reasoning-models#:~:Owkin...

[16] https://www.owkin.com/newsfeed/owkin-to-build-ai-agents-as-part-of-a-multi-year-k-pro-license-agreement-with-astrazeneca#:~:Owkin...

[17] https://www.mckinsey.com/industries/life-sciences/our-insights/reimagining-life-science-enterprises-with-agentic-ai#:~:AI%20...

[18] https://www.owkin.com/newsfeed/owkin-launches-k-pro-the-first-agentic-ai-co-pilot-for-biopharma-powered-by-biological-reasoning-models#:~:,%E2%...

[19] https://bio-in-tech.com/blog/ai-drug-discovery-market-size-growth-2026/#:~:Grand...

[20] https://www.owkin.com/newsfeed/owkin-launches-k-pro-the-first-agentic-ai-co-pilot-for-biopharma-powered-by-biological-reasoning-models#:~:At%20...

[21] https://www.owkin.com/newsfeed/owkin-to-build-ai-agents-as-part-of-a-multi-year-k-pro-license-agreement-with-astrazeneca#:~:Owkin...

[22] https://www.mckinsey.com/industries/life-sciences/our-insights/reimagining-life-science-enterprises-with-agentic-ai#:~:Gen%2...

[23] https://www.mckinsey.com/industries/life-sciences/our-insights/reimagining-life-science-enterprises-with-agentic-ai#:~:Lower...

[24] https://www.owkin.com/newsfeed/owkin-launches-k-pro-the-first-agentic-ai-co-pilot-for-biopharma-powered-by-biological-reasoning-models#:~:K%20P...

[25] https://aws.amazon.com/marketplace/pp/prodview-janhwocmshmxm#:~:AWS%2...

[26] https://www.businesswire.com/news/home/20260605704506/en/Owkin-to-Build-AI-Agents-as-Part-of-a-Multi-Year-K-Pro-Collaboration-with-Sanofi#:~:K%20P...

[27] https://bio-in-tech.com/blog/ai-drug-discovery-market-size-growth-2026/#:~:Key%2...

[28] https://www.mckinsey.com/industries/life-sciences/our-insights/reimagining-life-science-enterprises-with-agentic-ai#:~:disea...

[29] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:disco...

[30] https://www.mckinsey.com/industries/life-sciences/our-insights/reimagining-life-science-enterprises-with-agentic-ai#:~:delay...

[31] https://www.mobihealthnews.com/news/owkin-sanofi-expand-ai-collaboration-build-drug-development-ai-agents#:~:The%2...

[32] https://www.mobihealthnews.com/news/owkin-sanofi-expand-ai-collaboration-build-drug-development-ai-agents#:~:K%20P...

[33] https://www.owkin.com/newsfeed/owkin-to-build-ai-agents-as-part-of-a-multi-year-k-pro-license-agreement-with-astrazeneca#:~:Under...

[34] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:Daiic...

[35] https://www.businesswire.com/news/home/20260605704506/en/Owkin-to-Build-AI-Agents-as-Part-of-a-Multi-Year-K-Pro-Collaboration-with-Sanofi#:~:Durin...

[36] https://www.mobihealthnews.com/news/owkin-sanofi-expand-ai-collaboration-build-drug-development-ai-agents#:~:,Owki...

[37] https://www.owkin.com/newsfeed/owkin-to-build-ai-agents-as-part-of-a-multi-year-k-pro-license-agreement-with-astrazeneca#:~:,supp...

[38] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:Then%...

[39] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:to%20...

[40] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:Where...

[41] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:Immun...

[42] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:gener...

[43] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:Merck...

[44] https://hlth.com/insights/articles/q1-2026-when-pharma-ai-partnerships-crossed-a-threshold#:~:Tempu...

[45] https://www.genengnews.com/topics/drug-discovery/roche-genentech-innovation-center-boston-launched-at-harvards-enterprise-research-campus/#:~:hub%2...

[46] https://www.mobihealthnews.com/news/owkin-sanofi-expand-ai-collaboration-build-drug-development-ai-agents#:~:That%...

[47] https://www.mobihealthnews.com/news/owkin-sanofi-expand-ai-collaboration-build-drug-development-ai-agents#:~:In%20...

[48] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:,FDA%...

[49] https://apnews.com/article/004c0ce0442b72c37bfec6e032796808#:~:can%2...

[50] https://www.mobihealthnews.com/news/owkin-sanofi-expand-ai-collaboration-build-drug-development-ai-agents#:~:Under...

DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by Adrien Laurent, a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.