Evaluating Biotech AI Use Cases: A Prioritization Guide

6/26/2026 • 45 min read

biotech ai use cases | ai in drug discovery | biotech ai implementation framework | ai data readiness life sciences | pharma ai use case prioritization | biotech ai risk assessment | machine learning in biotech | artificial intelligence in biotechnology

Executive Summary

Biotechnology and pharmaceutical companies now face dozens of plausible artificial intelligence (AI) use cases, but very few organizations have a repeatable method for choosing which ones to fund. This report builds a prioritization framework for **biotech AI use cases** grounded in verified 2025 and 2026 data rather than vendor claims. The core finding is stark: a July 2025 MIT NANDA research initiative studying enterprise generative AI found that despite an estimated $30 billion to $40 billion in enterprise investment, **95% of organizations are getting zero return**, with only 5% of integrated pilots extracting measurable value ([1] mlq.ai). Deloitte's 2026 life sciences outlook corroborates this pattern within the industry specifically: 78% of biopharma and medtech executives expect AI to play a central role in driving major change in 2026, yet only 22% report having successfully scaled AI and just 9% report achieving significant returns ([2] deloitte.com).

The gap between adoption and value is not primarily a model-quality problem. It is a prioritization and data-readiness problem. A November 2025 survey of roughly 100 biotechnology and pharmaceutical organizations found the highest AI adoption in tasks with clean, verifiable data, such as literature review (76% adoption) and protein structure prediction (71% adoption), while only a small minority of surveyed organizations describe their R&D data as "fully integrated" across functions (a data-readiness gap detailed later in this report). Capgemini's global biopharma survey found that 82% of executives believe AI will fundamentally transform R&D and 79% are actively building integration strategies, yet the same organizations describe most AI efforts as still early-stage ([3] capgemini.com).

Real deployments illustrate what "crossing the divide" looks like in practice. Insilico Medicine initiated a Phase III trial in July 2026 for rentosertib, an AI-discovered and AI-designed oral TNIK inhibitor for idiopathic pulmonary fibrosis, whose Phase IIa results showed a mean forced vital capacity improvement of +98.4 mL at 12 weeks ([4] prnewswire.com). Bristol Myers Squibb's Devens, Massachusetts biologics site built more than 30 targeted machine learning use cases into its manufacturing workflow, cutting New Product Introduction time by 42% and lifting manufacturing volume by more than 40% ([5] aiformanufacturing.org). Takeda signed a deal worth up to $1.7 billion in milestone payments with Iambic Therapeutics for AI-driven small molecule discovery ([6] fiercebiotech.com), and Novo Nordisk deployed a governed reasoning agent on Microsoft Azure that reduced time-to-insight from weeks to minutes across more than 200,000 patient-years of harmonized clinical data ([7] microsoft.com).

Regulatory scaffolding is also maturing. The U.S. Food and Drug Administration (FDA) established a CDER (Center for Drug Evaluation and Research) AI Council in 2024 and, in January 2025, issued draft guidance proposing a seven-step, risk-based credibility assessment framework for AI models used to support regulatory decisions on drug safety, effectiveness, or quality ([8] fda.gov) ([9] dlapiper.com). This report walks through a practical prioritization methodology (impact, feasibility, regulatory complexity, and data readiness), quantifies where AI is already delivering return on investment (ROI) according to independent surveys from NVIDIA and Deloitte, catalogs the risks that most often sink pilots, and closes with implications for R&D, manufacturing, and commercial organizations planning AI investment through 2026 and beyond. Throughout, the analysis distinguishes vendor and consultancy claims from independently measured outcomes, and flags where public data is still thin.

Introduction and Background

Artificial intelligence is not new to biotechnology, but the volume and diversity of applications reaching production have expanded sharply. The FDA defines AI as "a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual

environments," with machine learning (ML) as the subset of techniques most commonly used across the drug product life cycle ([10] fda.gov). The agency reports "a significant increase in the number of drug application submissions using AI components" and notes these submissions span nonclinical, clinical, postmarketing, and manufacturing phases ([11] fda.gov), with CDER's own count reaching over 500 AI-component submissions between 2016 and 2023 ([12] fda.gov). The agency notes that AI is "increasingly integrated" into two areas where CDER is separately active: Digital Health Technologies (DHTs) and Real-World Data (RWD) analytics, both of which intersect with the clinical development use cases discussed later in this report ([13] fda.gov).

The market opportunity behind this activity is large but measured differently depending on scope. Coherent Market Insights estimates the global AI in life science market at **$4.68 billion** in 2026, growing to **$22.95 billion** by 2033 at a compound annual growth rate (CAGR) of **25.5%** ([14] coherentmarketinsights.com). A narrower segment, Precedence Research's AI in drug discovery market specifically, is sized at **$7.62 billion** in 2026, rising to **$17.81 billion** by 2035 at a CAGR of **9.90%** ([15] precedenceresearch.com). The discrepancy between the two figures illustrates a recurring problem when evaluating AI use cases: "AI in life sciences" is a much broader category (spanning diagnostics, manufacturing, patient monitoring, and regulatory affairs, not only discovery) than "AI in drug discovery" narrowly defined, and readers comparing market-sizing claims across sources should always check what activities are actually in scope. Software (algorithms, platforms, and modules) accounts for an estimated 56.3% share of the broader life science AI market in 2026 ([16] coherentmarketinsights.com). By deployment model, cloud-based infrastructure is expected to account for **76.7%** of the market in 2026, reflecting the scalability and collaboration advantages cloud platforms offer when processing the large, heterogeneous datasets common in life science research ([17] coherentmarketinsights.com).Regionally, North America is expected to hold a 40.3% share of this broader life science AI market in 2026, with Asia Pacific identified as the fastest-growing region at an estimated 23.4% share ([18] coherentmarketinsights.com).

Consolidation is also underway: in April 2025, Siemens AG acquired Dotmatics, a life sciences research and development (R&D) software company, for **$5.1 billion**, folding it into Siemens' AI-powered product lifecycle management portfolio ([19] coherentmarketinsights.com). This is symptomatic of a broader industry shift: rather than experimenting with isolated point tools, large pharmaceutical companies are increasingly signing multi-year platform and data-licensing agreements with AI-native biotechs, exemplified by Pfizer's 2026 licensing deal with Chai Discovery for generative antibody design ([20] intuitionlabs.ai) and Takeda's up to $1.7 billion agreement with Iambic Therapeutics ([6] fiercebiotech.com).

The pace of underlying model progress is also accelerating the range of use cases worth evaluating. A recent Nature paper describes a technique called neural iterative selection-expansion (NISE), which pairs a graph neural network for protein sequence design with a structure predictor in a closed iterative loop, using it "to design proteins that, using different folds, specifically bind to two chemically distinct small-molecule drugs, exatecan and apixaban, with success rates of 100% and 83%, respectively" ([21] nature.com). Separately, an npj Drug Discovery paper describes V-SYNTHES2, an upgraded workflow for structure-based virtual screening that expands coverage "to REAL Space of 36 billion readily available compounds," illustrating how computational screening capacity keeps expanding by orders of magnitude ([22] nature.com). The authors calculate that brute-force docking of the full 36-billion-compound library would require roughly 130 years of wall-clock time on a 320-CPU-core cluster at an estimated cloud cost of about $3.6 million, whereas their hierarchical decomposition approach reduces the required docking to under 4 million compounds at an estimated cost of about $380, a reduction of more than four orders of magnitude in compute cost (nature.com). Advances of this kind expand the technically feasible use case list faster than most organizations can absorb them, which is precisely why a prioritization discipline, rather than an inventory of interesting capabilities, is the more urgent need.

Yet enthusiasm and spend have consistently outpaced measured returns. This is the central tension this report addresses: given a long list of candidate use cases across discovery, clinical development, manufacturing, and commercial operations, how should a biotech or pharmaceutical organization decide what to build, what to buy,

and what to defer? The remainder of this report develops an evaluation framework, quantifies where AI is delivering measurable value today, surveys the governance and risk landscape that shapes what is feasible in regulated environments, and reviews five named deployments in detail before turning to implications for 2026 and beyond. Life sciences consultancies advising on this landscape, such as IntuitionLabs, which positions itself around AI and analytics solutions, Veeva ecosystem implementation, and advisory services for regulated pharmaceutical and life science organizations, frame the challenge similarly: durable AI value depends less on model selection than on data foundations, workflow integration, and governance discipline ([24] intuitionlabs.ai).

A Taxonomy of Biotech AI Use Cases

Before prioritizing use cases, organizations need a shared vocabulary for what "AI" actually covers, since the term spans a wide range of techniques with very different data requirements, risk profiles, and regulatory treatment.

Discovery and Design Use Cases

The Benchling 2026 Biotech AI Report, based on a November 2025 survey of roughly 100 biotechnology and pharmaceutical organizations actively using AI in R&D, breaks discovery and design use cases into distinct adoption tiers ([25] benchling.com):

- **Literature review and knowledge extraction**: 76% adoption, the highest of any category, because outputs are easy to verify against source documents.
- **Protein structure and property prediction**: 71% adoption, building on AlphaFold-class structure-prediction models.
- **Scientific reporting and internal communication**: 66% adoption.
- **Target identification**: 58% adoption today, with 18% of non-adopters planning to adopt.
- **Lead optimization and selection**: 47% adoption.
- **De novo generative drug design** and **synthetic biology and pathway design**: 42% adoption each, reflecting newer and more technically demanding categories.
- **Hit identification**: 37% adoption.
- **Biomarker identification and analysis**: 40% adoption (all figures per the same Benchling survey cited above).

Generative antibody design illustrates how quickly a discovery use case can move from academic curiosity to industrial deployment. Chai Discovery's Chai-2 model achieved a roughly **16% to 20% hit rate** (validated binding) in zero-shot design against previously unaddressed targets, more than 100-fold better than prior computational methods, which historically showed hit rates below 0.1% ([26] intuitionlabs.ai). Pfizer's June 2026 licensing agreement grants access to the successor Chai-3 model, described as doubling Chai-2's success rate, plus a custom model trained on Pfizer's proprietary data ([27] intuitionlabs.ai).

Experimentation, Data Science, and Portfolio Use Cases

Beyond discovery, AI is increasingly embedded in day-to-day experimentation and portfolio management. The same Benchling survey found **50% adoption** for data analytics and visualization, **39%** for experiment design and planning, **29%** for ADME (absorption, distribution, metabolism, and excretion) modeling, and **31%** for portfolio management and asset strategy ([28] benchling.com). Scientists have also shifted their default research

behavior: 89% now use AI copilots or reasoning tools as their first stop for interrogating and synthesizing data, ahead of manual search methods like Google or PubMed at 92% and internal non-AI dashboards at 71%. Adoption differs sharply between scientific and non-scientific workflows: 81% of respondents use AI for scientific use cases such as biomarker discovery or molecule design, while 92% use it for non-scientific tasks such as document authoring and software engineering (figures per the same survey).

Clinical Development Use Cases

In clinical development, AI use cases split between trial design and operational execution. Novo Nordisk's FounData initiative harmonized more than **200,000 patient-years** of clinical trial data across CDISC, SDTM, and ADaM standards, enabling a governed reasoning agent to reduce time-to-insight "from weeks to minutes" for many exploratory analyses (detailed further in the case studies section below). Academic research is pushing agentic architectures further: a peer-reviewed Nature Communications paper describes EmulatRx, a multi-agent framework that uses electronic health record (EHR) data to derive real-world evidence for clinical trial design, with specialized Trialist, Informatician, Clinician, and Statistician agents collaborating under a Supervisor agent ([29] nature.com). In one showcase, the system's Statistician agent calculated that a prospective trial replicating an acute heart failure study could be adequately powered with **3,107 patients** rather than the original cohort's **6,971**, using standard Cox proportional hazards sample-size methodology ([30] nature.com). The paper also reports that manual execution of comparable workflows, requiring domain experts to translate protocols and write analysis code, "typically spans multiple days to weeks," compared with median automated runtimes measured in single-digit to tens of minutes depending on the underlying large language model (LLM) ([31] nature.com).

Manufacturing and Supply Chain Use Cases

Manufacturing is where some of the most rigorously quantified biotech AI benefits have emerged, because process data is abundant and outcomes (yield, cycle time, deviations) are already tracked for quality purposes. Bristol Myers Squibb's Devens, Massachusetts site built a portfolio of **more than 30 distinct machine learning use cases**, applying predictive ML directly to fermentation parameters, cell culture conditions, and upstream and downstream process variables to anticipate deviations before they affected yield or quality ([32] aiformanufacturing.org). Separately, Nature Communications researchers demonstrated an integrated digital formulator and self-driving tableting data factory that reduced the time from material characterization to in-specification tablets to **6 hours** while cutting active pharmaceutical ingredient (API) material use by **65%** compared with prior state-of-the-art methods ([33] nature.com). The same paper's introduction cites estimates that AI-driven scientific advances "can shorten drug development timelines from ≈12 to 15 years to 3 to 4 years" as discovery and clinical bottlenecks are addressed, which is part of why the authors argue that chemistry, manufacturing, and controls (CMC) activities, not discovery, have become the next bottleneck to automate ([34] nature.com).

Consultancy Modus Create groups manufacturing and supply chain use cases into four recurring patterns across its life sciences client base: **deviation prediction**, where real-time monitoring detects anomalies before they affect production and reduces compliance risk; **process optimization**, where models continuously refine operating parameters to improve throughput, yield, and consistency; **batch record review**, where AI automates review of batch documentation to accelerate release cycles while maintaining audit readiness; and **predictive maintenance**, where equipment and sensor data forecast failures so interventions can be scheduled before downtime occurs ([35] moduscreate.com). These four patterns map closely onto the underlying logic of the BMS Devens portfolio: rather than one generalized manufacturing model, the highest-value deployments decompose the plant floor into discrete, independently validated prediction problems.

Commercial and Medical Affairs Use Cases

AI use cases are not confined to the laboratory or the factory floor; commercial and medical affairs functions represent a fourth major category. Modus Create identifies medical affairs and commercial operations as one of four life sciences value-chain areas "currently witnessing some of the highest returns from AI deployment," alongside research, clinical trials, and manufacturing ([36] moduscreate.com). Within this category, three use cases recur most often: **field insights and omnichannel intelligence**, which analyzes real-world engagement data to inform commercial strategy; **explainable medical chatbots**, which provide transparent, scientifically validated guidance to healthcare professionals (HCPs) and internal teams; and **content generation and MLR (medical, legal, and regulatory) acceleration**, which automates creation and review of promotional and medical content to improve speed and consistency while maintaining compliance ([37] moduscreate.com). This category is instructive for the prioritization framework developed later in this report because it tends to carry lower regulatory complexity than clinical or manufacturing use cases (most commercial AI does not directly inform a regulatory submission) while still requiring rigorous MLR review workflows, illustrating that "low regulatory complexity" is relative, not absolute, and still demands compliance-aware design.

Building a Prioritization Framework for Biotech AI Use Cases

With dozens of candidate use cases, biotechs need a repeatable scoring method rather than ad hoc enthusiasm for the newest model release. Consultancy Modus Create recommends prioritizing use cases "based on potential impact, feasibility, regulatory complexity, and resource availability," running tightly scoped proofs of concept (POCs) with clear success criteria, and defining quantitative and qualitative key performance indicators (KPIs) before each project begins ([38] moduscreate.com). ZS Associates, in a white paper co-authored with AstraZeneca, notes that pharma R&D AI investments "have already yielded impressive results such as faster target identification, accelerated market research insights, improved clinical trial recruitment rates and significant cost savings" ([39] zs.com), while cautioning that "current ROI metrics often focus on isolated operational outputs like time or cost savings" and that these are frequently "finger-in-the-air estimates" rather than rigorous value quantification, proposing instead a multidimensional ROI framework spanning scientific, operational, and commercial dimensions ([40] zs.com).

Dimension One: Business Impact

The first screening dimension is potential business impact, measured against concrete KPI categories. Modus Create's recommended KPI framework spans three tiers: **business** metrics (time-to-market reduction, cost savings, revenue impact), **operational** metrics (process cycle time, deviation reduction, trial patient dropout), and **data and AI** metrics (model accuracy, drift detection, data quality). Deloitte's 2026 survey found that 41% of biopharma executives cite improving R&D productivity as their top cost-management priority, reflecting the fact that the average cost of bringing a new drug to market now tops **$2 billion** ([41] deloitte.com), while 47% of medtech leaders name AI implementation as their primary cost-containment strategy for 2026 ([42] deloitte.com). This cost pressure compounds with competitive pressure: 37% of surveyed executives cited increased competition from generics and biosimilars as a top 2026 trend, and 26% cited the expiration of blockbuster drugs and impending patent cliffs, both of which raise the stakes for any AI investment that can shorten development timelines or extend a product's competitive window ([43] deloitte.com).

Dimension Two: Technical and Data Feasibility

The second dimension is whether the organization's data foundation can actually support the use case. This is where most biotech AI initiatives fail. Benchling's survey found that only 6% of organizations describe their R&D data integration as "fully integrated," with 45% describing it as merely "developing" (key datasets connected, but inconsistent standards and gaps remain) ([44] benchling.com). Ardigen's 2026 industry retrospective makes the point directly: "the limiting factor is rarely access to algorithms. It is the ability to transform heterogeneous biological data into systems that produce reliable, interpretable insight over time" ([45] ardigen.com). The retrospective adds that "integration, quality control, and reproducibility often decide whether an AI component becomes an enduring capability or stays a one-off success," a distinction directly relevant to the data-feasibility screen in this framework ([46] ardigen.com). Modus Create similarly warns against "starting without data governance," recommending organizations build a data foundation for compliant AI before scaling any use case (a theme developed further below).

Dimension Three: Regulatory Complexity

The third dimension, regulatory complexity, separates use cases that touch patient safety, product quality, or regulatory submissions from those that do not. The FDA's January 2025 draft guidance explicitly carves out two categories that fall outside its risk-based credibility assessment framework: AI used in drug discovery, and AI used to streamline internal operations (such as drafting a regulatory submission) that does not affect patient safety, drug quality, or the reliability of clinical or nonclinical study results ([47] dlapiper.com). Everything else, including AI supporting trial design, safety signal detection, and manufacturing quality decisions, falls within scope and requires a seven-step credibility assessment: define the question of interest, establish the model's context of use (COU), assess model risk, develop a credibility assessment plan, execute it, document results, and determine adequacy ([9] dlapiper.com). If a sponsor cannot adequately establish an AI model's credibility for its intended context of use, the draft guidance outlines several remedial paths, including introducing additional supplementary evidence, increasing the rigor of credibility assessment activities, implementing risk-mitigating controls, or modifying the modeling approach altogether ([48] dlapiper.com). Use cases with low regulatory complexity (internal literature review, scientific reporting) can move to production quickly; those with high complexity (AI-informed dosing decisions, automated batch release) require substantially more validation investment before any ROI materializes.

Dimension Four: Resource Availability and Organizational Readiness

The fourth dimension considers whether the organization actually has the talent, infrastructure, and executive sponsorship to sustain a use case past the pilot stage. Benchling found that internal upskilling of existing scientific staff is the leading source of AI talent (67%), well ahead of hiring from technology companies (21%), and that AI leadership structures remain fragmented: 30% of surveyed organizations locate AI leadership within R&D, 24% within IT, 22% describe it as distributed, and only 16% have a dedicated Chief AI Officer or VP of AI and Data Science ([49] benchling.com). Deloitte found that only 14% of surveyed life sciences executives report full implementation of AI into daily workflows, with another 40% still working toward that goal ([50] deloitte.com).

Table 1 below summarizes how these four dimensions apply across representative use case categories, drawing on the adoption and impact figures cited throughout this report.

Use Case Category	Typical Adoption (2025 to 2026)	Data Feasibility	Regulatory Complexity	Representative Impact Evidence
Literature review and knowledge extraction	76% (per Benchling, cited above)	High (structured text corpora)	Low (outside FDA's credibility-assessment scope for most uses)	Fastest-adopted "killer app" among surveyed biotechs
Protein structure and property prediction	71% (per Benchling, cited above)	High (large public structural datasets)	Low to moderate	Zero-shot antibody design at 16 to 20% hit rate, over 100x prior methods ([26] intuitionlabs.ai)
Target identification	58% (per Benchling, cited above)	Moderate (multi-omics integration required)	Low (falls under drug discovery carve-out) ([47] dlapiper.com)	50% of surveyed biotechs report faster time-to-target today
Clinical trial design and RWE analysis	Emerging, agentic frameworks in peer review ([29] nature.com)	Moderate to high, depends on EHR/CDISC harmonization	High (informs regulatory submissions)	Sample-size reduction from 6,971 to 3,107 patients in one showcase ([30] nature.com)
Biologics manufacturing process control	Portfolio approach (30+ use cases at one site) ([32] aiformanufacturing.org)	High (abundant sensor and batch data)	High (GxP, quality risk management)	42% NPI time reduction, 40%+ volume increase ([5] aiformanufacturing.org)

The table shows a consistent pattern: use cases with abundant, well-structured data and low regulatory exposure (literature review, structure prediction) reach high adoption fastest, while use cases with strong potential impact but higher data or regulatory complexity (clinical trial design, manufacturing quality) require more deliberate investment before returns appear, but can also deliver the largest quantified gains once mature, as the Devens and EmulatRx examples illustrate.

Implementation Considerations and Process Changes

Selecting the right use case is necessary but not sufficient. Organizations that successfully scale biotech AI initiatives make several structural changes to how they run pilots, govern data, and manage risk.

Data Governance as a Prerequisite, Not an Afterthought

Capgemini's Research Institute survey found that 82% of biopharma executives believe AI will fundamentally transform R&D, and 63% anticipate that most new molecular entities (NMEs) will originate from AI-driven platforms within the next decade ([51] capgemini.com). Yet the same survey's framing acknowledges that "to fully harness AI's potential in R&D, biopharma organizations must be data-ready," and that most organizations still lag in this dimension ([52] capgemini.com). Novo Nordisk's approach illustrates what readiness looks like in practice: years spent building the FounData platform on open industry standards (CDISC, SDTM, ADaM) before layering an AI agent on top meant that analyses could run on already-prepared data without bespoke pre-processing, as detailed in the case studies section below.

Governance, Validation, and MLOps for GxP Environments

For use cases that touch Good "x" Practice (GxP) regulated processes, formal AI governance is required, not optional. ISPE's (International Society for Pharmaceutical Engineering) guidance describes AI governance as "a system of laws, regulations, policies, controls, frameworks, standards… to manage, regulate, and optimize… AI technology," recommending alignment with ISO/IEC Standard 42001 and integration with the existing GAMP 5 (Good Automated Manufacturing Practice) risk-based approach to compliant GxP computerized systems ([53] ispe.org). Quality risk management processes under ICH (International Council for Harmonisation) Harmonized Tripartite Guideline Q9 must be applied, with the overall risk level for any AI system determined by its potential impact on patient safety, product quality, and data integrity ([54] ispe.org). ISPE's governance policy checklist further calls for a data privacy and protection policy that ensures "compliance with data protection laws such as… the General Data Protection Regulation, California Consumer Privacy Act, and EU AI Act," alongside separate policies for fairness and nondiscrimination, transparency and explainability, and stakeholder engagement ([55] ispe.org). Machine learning operations (MLOps) practices, including version control, continuous monitoring, model retraining triggers, and rollback and disaster recovery procedures, sit alongside AI governance as a parallel discipline needed to keep validated models compliant as they drift over time.

The pace of academic and preprint research on agentic discovery tools illustrates how quickly the underlying capability set is expanding. A 2026 preprint describes Rhizome OS-1, a semi-autonomous operating system in which multi-modal AI agents function as computational chemists, medicinal chemists, and patent agents, writing and executing their own analysis code and adapting generation strategies based on empirical screening feedback ([56] arxiv.org). In two oncology campaigns, the agent team executed 26 design seeds and produced 5,231 molecules, of which 91.9% of generated scaffolds were absent from the ChEMBL reference database of known bioactive compounds, indicating genuinely novel chemical matter rather than restatements of previously catalogued molecules ([57] arxiv.org). Because this is a preprint that has not yet completed peer review, its findings should be treated as preliminary relative to the peer-reviewed Nature and Nature Communications papers cited elsewhere in this report.

Running Controlled Pilots with Defined Success Criteria

Modus Create's guidance for scaling AI in life sciences recommends starting with "a high-value, high-feasibility use case" rather than a moonshot project, running "tightly scoped POCs with clear success criteria and active involvement from domain experts," and avoiding the common failure mode of "running too many unscalable POCs" that never graduate to production ([58] moduscreate.com). NVIDIA's second annual healthcare AI survey found that among healthcare and life sciences organizations broadly, 85% of executives said their AI budgets will increase in the coming year, and for almost half (46%) that increase will exceed 10% ([59] blogs.nvidia.com), suggesting that budget is rarely the binding constraint; well-governed pilot design usually is.

Scaling with MLOps and Continuous Improvement

Once a pilot demonstrates value, the harder problem is keeping it valid as data distributions, assay protocols, and manufacturing conditions drift over time. ISPE's GxP governance guidance treats MLOps as a distinct discipline running parallel to AI governance, encompassing version control, continuous integration and deployment, model monitoring and logging, and defined model retraining and fine-tuning triggers ([60] ispe.org). A risk-based validation review published in a peer-reviewed biosciences journal similarly emphasizes that "regulatory guidance from the FDA (2025) emphasizes that AI/ML systems must include predetermined retraining triggers, continuous monitoring pipelines, and documented risk mitigation strategies, particularly in contexts where clinical or quality-critical decisions are made" ([61] biotech-asia.org). The same review notes that

ISO/IEC 42001, published in 2023, "introduces the first AI Management System Standard (AIMS), focusing on lifecycle governance, transparent documentation, algorithmic fairness, and risk-based control mechanisms," and has been adopted globally with accredited certifications issued as recently as mid-2025 ([62] biotech-asia.org). Modus Create frames the practical version of this discipline as building continuous improvement loops once a use case scales, so that model performance regressions are caught operationally rather than discovered downstream in a quality deviation.

Cross-Functional Ownership and Talent Structures

The Deloitte survey found agency restructuring, tariffs, and other economic policy changes are competing for executive attention alongside AI investment (36% and 39% of respondents respectively cite these as top 2026 concerns) ([63] deloitte.com), meaning AI initiatives must compete for organizational bandwidth against other transformation priorities. Effective governance structures assign explicit accountability across functions: IT and security teams support scalable, secure infrastructure; compliance and quality assurance (QA) embed GxP considerations early; and business leaders drive adoption and prioritize which use cases advance.

Data Analysis and Evidence

This section consolidates the quantitative evidence on biotech AI adoption, investment, and return across the independent surveys reviewed for this report. *Table 2* below summarizes the headline figures from the four largest independent surveys cited in this report side by side, since each covers a different population and is easy to conflate without a direct comparison.

Survey	Sample	AI Adoption	Reported ROI / Scaling
NVIDIA State of AI in Healthcare and Life Sciences (2026)	Healthcare, pharma, biotech, medtech, payer/provider executives	70% actively using AI, up from 63% in 2024 (cited above)	85% say AI increases revenue, 80% say it reduces costs (cited above)
Deloitte 2026 Life Sciences Outlook	Global biopharma and medtech executives	48% cite accelerated digital transformation as an influential 2026 trend ([64] deloitte.com)	Only 22% have successfully scaled AI; 9% report significant returns ([2] deloitte.com)
Benchling 2026 Biotech AI Report	~100 self-selected AI-adopting biotech and biopharma organizations	76% adoption for literature review, the highest single use case (cited above)	50% report faster time-to-target; only 6% describe R&D data as fully integrated
MIT NANDA State of AI in Business (2025)	300+ disclosed AI initiatives, 52 organizations, 153 leaders (cross-industry, not life-sciences-specific)	$30 billion to $40 billion in enterprise generative AI investment tracked ([1] mlq.ai)	95% of organizations report zero measurable P&L return; just 5% of integrated pilots extract real value ([65] mlq.ai)

Read together, the table shows that adoption figures vary considerably depending on who is sampled: Benchling's self-selected AI-adopting biotechs naturally show higher use-case-level adoption than NVIDIA's broader healthcare sample or Deloitte's general biopharma and medtech executive base, while the cross-industry MIT NANDA data suggests the gap between piloting AI and capturing measurable value is not unique to life sciences but may be even more pronounced within it, given Deloitte's 9% figure sits below MIT NANDA's already-low 5% of pilots extracting real value once general enterprise scaling challenges are compounded by regulatory and data-integration complexity specific to biotech.

Adoption is rising across the healthcare and life sciences sector broadly. NVIDIA's 2026 survey found 70% of respondents said their organizations are actively using AI, up from 63% in 2024, and 69% reported using

generative AI and large language models, up from 54% ([66] blogs.nvidia.com). Within the pharmaceutical and biotechnology segment specifically, 57% of respondents said drug discovery is being driven by AI, and 46% named drug discovery and development as among their top ROI use cases ([67] blogs.nvidia.com). Across the full survey sample, 85% of executives said AI is helping increase revenue and 80% said it is helping reduce costs; 47% reported using or assessing agentic AI systems that act autonomously to achieve defined goals; and 82% said open-source software and models are moderately to extremely important to their AI strategy ([68] blogs.nvidia.com).

Deloitte's 2026 life sciences outlook, drawing on a separate global survey of biopharma and medtech executives, quantifies the maturity gap directly: while 48% of respondents identified accelerated digital transformation as a trend likely to substantially affect their organization in 2026, and 41% cited generative AI proliferation as an influential trend, only **22%** of life sciences leaders said they have successfully scaled AI and just **9%** reported achieving significant returns on those efforts ([64] deloitte.com). Agentic AI was a new category in this year's survey, cited by 30% of respondents as an influential trend ([69] deloitte.com).

Benchling's biotech-specific survey (fielded among self-selected AI-adopting organizations, not a general industry sample, and therefore likely to show higher adoption than the industry average) reported that among R&D tech budgets, the share allocated to AI clusters below 20% for most respondents, and that trust in raw LLM outputs remains uneven across scientific versus operational tasks. The survey's respondent base skewed toward larger organizations (53% with more than 1,000 employees) and toward North America (70%) ([70] benchling.com), a sampling detail worth noting when generalizing its findings to smaller biotechs or non-U.S./European markets.

The originating MIT NANDA research, based on a systematic review of over 300 publicly disclosed AI initiatives, structured interviews with 52 organizations, and survey responses from 153 senior leaders, found the divide between AI adopters and AI value-capturers is "determined by approach" rather than model quality or regulation, and that this pattern held "across both buyers (enterprises, mid-market, SMBs) and builders (startups, vendors, consultancies)" ([71] mlq.ai). The report frames the split starkly: "just 5% of integrated AI pilots are extracting millions in value, while the vast majority remain stuck with no measurable P&L impact" ([65] mlq.ai). The methodology behind these findings combined "a systematic review of over 300 publicly disclosed AI initiatives, structured interviews with representatives from 52 organizations, and survey responses from 153 senior leaders collected across four major industry conferences" ([72] mlq.ai). This finding is not specific to life sciences, but Ardigen's biotech-focused 2026 retrospective explicitly cites it as the framing context for why biotech AI conversations shifted in 2025 "away from isolated model performance toward system-level questions" ([73] ardigen.com).

Cybersecurity concerns are rising in step with AI adoption rather than being displaced by it: Deloitte found 35% of surveyed life sciences executives now name cybersecurity as a factor likely to affect their 2026 strategy, alongside the 41% citing generative AI proliferation and 30% citing agentic AI as influential trends ([74] deloitte.com). This aligns with the governance emphasis of the ISPE and FDA frameworks discussed earlier in this report, which direct organizations to "implement a policy to ensure the security and safety of AI systems" and "protect AI systems from unauthorized access, tampering, and attacks" as a foundational governance policy rather than an optional add-on ([75] ispe.org).

On regulatory data, the FDA's own count shows the scale of submitted evidence: over 500 submissions containing AI components between 2016 and 2023, plus more than 800 external comments received on a May 2023 discussion paper, informing the January 2025 draft guidance ([76] fda.gov). On market growth, the two market-sizing estimates cited earlier in this report (Coherent Market Insights' broader $4.68 billion to $22.95 billion life science AI market at 25.5% CAGR, versus Precedence Research's narrower $7.62 billion to $17.81 billion drug discovery AI market at 9.90% CAGR) both point to durable double-digit growth, even though the narrower drug-discovery segment's CAGR is meaningfully lower than the broader category's, likely reflecting the larger base of adjacent AI applications (diagnostics, manufacturing, patient monitoring) captured in the wider

definition ([14] coherentmarketinsights.com) ([15] precedenceresearch.com). Precedence Research further breaks down the drug-discovery segment geographically and by application: North America accounted for the largest market share, 56.18%, in 2025, while the Asia-Pacific region is projected to grow at a CAGR of 21.1% from 2026 to 2035, and the oncology therapeutic area alone accounted for 21% of 2025 revenue ([77] precedenceresearch.com).

Regional and sector-level survey data reinforce the same optimism-versus-execution gap documented above. Deloitte found that 90% of biopharma leaders in surveyed European and Asian countries report positive or cautiously positive expectations for 2026, with 83% predicting steady or strong revenue growth, while only 56% of US-based biopharma leaders share that positive outlook and 27% remain negative or uncertain ([78] deloitte.com). Medtech executives were more consistently optimistic across geographies, with 84% of US medtech leaders and 78% of non-US medtech leaders anticipating similar revenue growth of roughly 81% overall ([79] deloitte.com). Capgemini's Research Institute adds a discovery-specific data point: 74% of biopharma executives believe generative AI holds significant potential specifically in drug discovery, a narrower and more bullish claim than the general R&D transformation figures cited earlier ([80] capgemini.com).

Case Studies and Real-World Examples

Insilico Medicine: Rentosertib Reaches Phase III

Insilico Medicine (HKEX: 3696) announced in July 2026 the initiation of a Phase III clinical trial for **rentosertib** (formerly ISM001-055/INS018_055), a potentially first-in-class oral small-molecule inhibitor of TNIK (TRAF2- and NCK-interacting kinase) for idiopathic pulmonary fibrosis (IPF), a progressive, age-related fibrotic lung disease ([81] prnewswire.com). The molecule was discovered and designed through Insilico's Pharma.AI platform, combining a fibrosis target prioritized by the company's Biology42 (PandaOmics) AI biology engine with a small molecule generated by its Chemistry42 AI chemistry engine ([82] prnewswire.com). The GENESIS-IPF Phase IIa study reported manageable safety and tolerability, with the 60 mg once-daily dosing arm showing a mean forced vital capacity improvement of **+98.4 mL** at 12 weeks ([4] prnewswire.com). Insilico CEO Alex Zhavoronkov described the milestone as "a major milestone for the company and my life as well," adding that if the drug outperforms standard of care with a disease-modifying effect, "it's going to be the first and only of its kind" ([83] bio-itworld.com). Trade press covering the milestone described it as "the next step in making AI-discovered drugs a reality for patients," underscoring that the Phase III initiation, not the earlier discovery milestones, is what the industry is treating as the real proof point ([84] bio-itworld.com). This case demonstrates AI use spanning the full discovery-to-clinic pipeline (target discovery, molecule design, and biology-driven indication selection), with the discovery-to-clinic journey separately published in Nature Biotechnology and the Phase IIa results in Nature Medicine ([85] prnewswire.com).

Bristol Myers Squibb: Portfolio-Scale Manufacturing AI at Devens

Bristol Myers Squibb's Devens, Massachusetts facility, one of its most advanced biologics and cell therapy manufacturing sites, faced a persistent challenge: biological variability in living-cell-based production slowed New Product Introduction (NPI) cycles and constrained throughput ([86] aiformanufacturing.org). Rather than deploying a single platform, the site built a portfolio of more than 30 distinct machine learning use cases, applying predictive ML directly to fermentation parameters, cell culture conditions, and upstream and downstream process variables ([32] aiformanufacturing.org). The results were measured across three dimensions

simultaneously: a **42%** reduction in NPI time, a **40%+** increase in manufacturing volume without proportional capital investment, and a **40%+** reduction in emissions from efficiency gains in energy and resource use ([5] aiformanufacturing.org). The site was subsequently recognized by the World Economic Forum's Global Lighthouse Network in January 2026 ([87] aiformanufacturing.org). A key takeaway the case study authors emphasize is that "domain expertise must lead AI deployment": BMS paired process scientists with data teams so that models reflected actual biological mechanisms rather than statistical artifacts ([88] aiformanufacturing.org). A second takeaway the authors highlight is that "sustainability gains follow efficiency gains," meaning the emissions reduction at Devens "emerged as a byproduct of optimized resource use, not a separate initiative" ([89] aiformanufacturing.org).

Takeda and Iambic Therapeutics: A Milestone-Based AI Partnership

In February 2026, Takeda Pharmaceuticals signed a multiyear technology and discovery agreement with Iambic Therapeutics for access to Iambic's suite of AI-driven drug discovery platforms, including its generative model NeuralPLexer, designed to predict protein-ligand complexes ([90] fiercebiotech.com). The deal covers upfront research-cost and technology-access payments, plus milestone-based payments that could exceed **$1.7 billion**, applied to small molecule programs in oncology, gastrointestinal, and inflammation therapeutic areas ([6] fiercebiotech.com). Iambic CEO Tom Miller framed the deal as evidence the technology has moved past speculative hype: "As you think about the AI space, work your way through landscape, and categorize Iambic as company, it is not about AI hype. We are showing in patients that these drugs are better, creating validation" ([91] fiercebiotech.com). Notably, Iambic's most advanced clinical-stage asset, oral oncology candidate IAM1363, delivered promising early-phase clinical data prior to the company's own $100 million funding round in November 2025, underscoring that milestone-based biobucks deals of this kind reward demonstrated clinical translation, not model benchmarks alone ([92] fiercebiotech.com).

Novo Nordisk: Governed Reasoning Agents on Azure

Novo Nordisk built its FounData initiative to harmonize more than 200,000 patient-years of clinical trial data across studies, disease areas, and research programs, aligned to open industry standards including CDISC, SDTM, and ADaM ([7] microsoft.com). Working with Microsoft's AI Acceleration Studio, the company then deployed a governed reasoning agent that allows researchers to analyze clinical data while maintaining rigor, oversight, and compliance. Before production rollout, teams executed thousands of automated tests and implemented layered evaluation combining query validation, system performance metrics, expert review, and continuous user feedback, alongside human-in-the-loop workflows where biostatisticians validate promising outputs before they influence major decisions ([93] microsoft.com). FounData's Rasmus S. Andersen summarized the governance philosophy: "The agent is there to inform decisions, not to take over control. It drafts the analysis; the scientist decides what to do with it. That distinction is foundational to how we built it" ([94] microsoft.com). The reported business impact: the number of scientific ideas the team can evaluate per quarter rose from a capacity of "5 to 10 strong ideas" to more than 50 ([95] microsoft.com).

(Hypothetical Example) A Mid-Size Biotech Prioritizing Its First Three AI Use Cases

Consider a hypothetical 400-person clinical-stage biotech deciding where to invest its first dedicated AI budget. Applying the four-dimension framework from this report, the company might screen a list of ten candidate use cases (ranging from a literature-review copilot to an agentic clinical trial design assistant) and find that the literature-review copilot scores highest on data feasibility (structured public corpora) and lowest on regulatory complexity (falls outside the FDA's credibility-assessment scope), making it a reasonable first deployment despite modest headline impact. A generative lead-optimization tool might score highest on potential impact but lowest on data feasibility, since the company's own assay data is fragmented across three electronic lab notebook systems, implying a data-integration project should precede (not accompany) that AI investment. This hypothetical sequencing, low regulatory complexity and high data feasibility first, higher-impact but higher-complexity use cases second, mirrors the pattern observed across the four real-world cases above, where the fastest-adopted use cases (literature review, structure prediction) also had the cleanest underlying data.

Implications and Future Directions

Several trends emerging from this evidence base will shape how biotech AI use cases evolve through the remainder of 2026 and beyond. First, agentic AI, systems that autonomously plan and execute multi-step tasks rather than responding to single prompts, is moving from a survey afterthought to a tracked category: Deloitte added it as a new metric this year and found 30% of life sciences respondents already cite it as an influential trend ([69] deloitte.com), while NVIDIA's broader healthcare survey found 47% already using or assessing agentic systems (cited in the data analysis section above). The EmulatRx clinical trial design framework and Novo Nordisk's governed reasoning agent (profiled in the case studies section) both illustrate what disciplined agentic deployment looks like in a regulated setting: multiple specialized agents with clear role boundaries, human-in-the-loop validation gates, and explicit logging for auditability ([29] nature.com).

Second, regulatory frameworks will continue to formalize around risk-based tiers rather than blanket rules. The FDA's CDER AI Council, established in 2024, is explicitly tasked with promoting "consistency on CDER considerations related to AI when evaluating drug safety, effectiveness, and quality" as submissions volume continues to rise ([96] fda.gov). Biotechs planning AI-informed regulatory submissions should expect to engage the agency early using the context-of-use and credibility-assessment vocabulary the guidance establishes, since the level of evidence required scales directly with model risk ([97] dlapiper.com).

Third, the divide between AI adopters and AI value-capturers documented at the enterprise level by MIT NANDA is unlikely to close through better models alone; it requires organizational changes to data governance, cross-functional ownership, and pilot design ([71] mlq.ai). Life sciences consultancies advising on this transition, including firms like IntuitionLabs that specialize in Veeva ecosystem implementation and pharmaceutical AI and analytics advisory rather than in selling a competing foundation model, tend to frame their value proposition around exactly this integration and governance layer, positioning themselves as complementary to, rather than a substitute for, the underlying AI platforms and vendors an organization selects ([98] intuitionlabs.ai). This complementary-advisor posture reflects a broader market structure: platform vendors (Insilico, Chai Discovery, Iambic) compete on model capability, cloud hyperscalers (Microsoft, NVIDIA) compete on infrastructure, and consultancies compete on implementation and governance discipline, three distinct roles that biotech buyers should evaluate separately rather than expecting a single vendor to fill all three.

Fourth, the "build versus partner" decision will keep shaping deal structures across the industry. Ardigen's 2026 retrospective observed that 2025 already "reinforced the move toward platform-oriented strategies," with large pharma companies "increasingly pursuing longer-term approaches that integrate data, models, and workflows rather than deploying single task tools" ([99] ardigen.com). Takeda's pattern of stacking multiple milestone-based AI partnerships (Iambic Therapeutics for up to $1.7 billion, alongside a separate October 2026 AI partnership with drug designer Nabla Bio) shows that large pharmaceutical companies increasingly treat external AI-native

biotechs as an extension of their own discovery engine rather than a one-off vendor relationship ([100] fiercebiotech.com). The Nabla-Takeda deal could reach more than $1 billion in biobucks if future milestones are met, and Nabla's Joint Atomic Model is intended to generate new antibodies and multispecifics across Takeda's early-stage programs ([101] fiercebiotech.com). Pfizer's Chai Discovery deal follows the same logic, pairing early access to a frontier model with a custom pipeline trained on Pfizer's own proprietary data so that the partnership compounds over time rather than expiring with a single licensing fee ([102] intuitionlabs.ai). Smaller biotechs without the balance sheet for billion-dollar biobucks deals should expect the same underlying prioritization logic to apply at a smaller scale: partner for capability where a mature external platform already demonstrates clinical translation, and build in-house only where the resulting model or dataset is a genuine, defensible source of competitive advantage.

Fifth, talent and organizational structure will likely remain the binding constraint longer than compute or model access. With only 16% of surveyed biotechs reporting a dedicated Chief AI Officer or VP-level AI leader and internal upskilling remaining the dominant talent source (both figures cited earlier in this report), organizations that invest early in structured AI leadership and cross-functional governance are likely to compound their advantage over slower-moving peers, consistent with Deloitte's finding that "surveyed executives from organizations that are further along the AI maturity continuum express greater optimism about the economy, the industry, and their own financial performance" ([103] deloitte.com).

Frequently Asked Questions (FAQs)

What are the most common AI use cases in biotech today? The highest-adoption use cases, per Benchling's 2026 survey (cited in the taxonomy section above), are literature review and knowledge extraction (76% adoption), protein structure and property prediction (71%), scientific reporting (66%), and target identification (58%).

How should a biotech prioritize which AI use cases to fund first? Consultancy guidance converges on a four-dimension approach (business impact, technical and data feasibility, regulatory complexity, and resource or organizational readiness), with tightly scoped pilots and predefined success criteria before scaling, per the prioritization framework detailed earlier in this report ([40] zs.com).

Is AI data readiness really the main bottleneck, rather than model capability? Available evidence supports this. Only 6% of AI-adopting biotechs describe their R&D data as fully integrated (per Benchling's survey, cited in the prioritization framework section above), and both Ardigen and Capgemini independently identify data readiness, not algorithm access, as the limiting factor for durable value ([45] ardigen.com) ([52] capgemini.com).

Which biotech AI use cases carry the highest regulatory risk? Use cases that generate information intended to support regulatory decisions on drug safety, effectiveness, or quality, such as AI-informed clinical trial design or manufacturing quality decisions, fall within the FDA's risk-based credibility assessment framework and require the most extensive validation. Drug discovery and internal-operations AI (not affecting patient safety, quality, or study reliability) are explicitly carved out of this framework ([47] dlapiper.com).

How big is the AI in biotech and life sciences market? Estimates vary by scope. Coherent Market Insights sizes the broader AI in life science market at $4.68 billion in 2026, growing to $22.95 billion by 2033 (25.5% CAGR) ([14] coherentmarketinsights.com), while Precedence Research sizes the narrower AI in drug discovery market at $7.62 billion in 2026, growing to $17.81 billion by 2035 (9.90% CAGR) ([15] precedenceresearch.com).

Do most biotech AI pilots actually deliver ROI? Not yet, at the enterprise level generally: 95% of organizations studied by MIT NANDA saw zero measurable return from generative AI pilots ([104] mlq.ai). Within life sciences specifically, Deloitte found only 9% of surveyed executives report achieving significant AI returns, even though 78% expect AI to play a central role in 2026 ([2] deloitte.com). However, specific well-governed deployments,

such as BMS's Devens manufacturing site and Novo Nordisk's FounData agent, do show substantial, independently reported gains ([5] aiformanufacturing.org) (see case studies section).

Conclusion

Biotech AI use cases are no longer scarce; if anything, the challenge in 2026 is the opposite of scarcity, an overabundance of plausible candidates competing for the same finite budget, data engineering capacity, and regulatory attention. The evidence assembled in this report points to a clear pattern: use cases that succeed quickly share abundant, well-structured underlying data and low regulatory exposure, such as literature review and protein structure prediction, while use cases with the largest potential impact, including AI-driven clinical trial design and manufacturing quality control, require more deliberate investment in data governance and risk-based validation before returns materialize, but can then deliver outsized, well-documented gains, as the Bristol Myers Squibb Devens and Novo Nordisk FounData cases demonstrate.

The prioritization framework developed here (business impact, data feasibility, regulatory complexity, and organizational readiness) gives biotech and pharmaceutical leaders a structured alternative to funding whichever use case generated the most recent conference buzz. Applied consistently, it directs early investment toward high-feasibility, low-complexity use cases that build organizational AI fluency and data infrastructure, while reserving higher-risk, higher-impact use cases for later stages once governance maturity and data integration can support them. Given that independent surveys converge on the same finding, that adoption is outpacing measured value across the industry, the organizations most likely to close that gap in 2026 and beyond will be those that treat prioritization, data readiness, and governance as the primary constraints on biotech AI success, not model access or compute budget.

External Sources

[1] https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf#:~:Despi...

[2] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Despi...

[3] https://www.capgemini.com/wp-content/uploads/2026/01/120126CRI_Gen-AI-in-Lifesciences-_Final-interactive.pdf#:~:Our%2...

[4] https://www.prnewswire.com/news-releases/insilico-initiates-phase-iii-clinical-trial-for-rentosertib-its-ai-empowered-tnik-inhibitor-for-idiopathic-pulmonary-fibrosis-302819553.html#:~:The%2...

[5] https://aiformanufacturing.org/case-studies/bristol-myers-squibb-cuts-npi-time-42-and-boosts-biologics-volume-40-with-ai-at-devens#:~:42%25...

[6] https://www.fiercebiotech.com/biotech/takeda-inks-17b-ai-discovery-deal-iambic-therapeutics#:~:After...

[7] https://www.microsoft.com/en/customers/story/26569-novo-nordisk-as-azure#:~:The%2...

[8] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:The%2...

[9] https://www.dlapiper.com/en-us/insights/publications/2025/01/fda-releases-draft-guidance-on-use-of-ai#:~:FDA%2...

[10] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:Artif...

[11] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:FDA%2...

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

[13] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:Addit...

[14] https://www.coherentmarketinsights.com/industry-reports/artificial-intelligence-in-life-science-market#:~:The%2...

[15] https://www.precedenceresearch.com/artificial-intelligence-in-drug-discovery-market#:~:The%2...

[16] https://www.coherentmarketinsights.com/industry-reports/artificial-intelligence-in-life-science-market#:~:The%2...

[17] https://www.coherentmarketinsights.com/industry-reports/artificial-intelligence-in-life-science-market#:~:The%2...

[18] https://www.coherentmarketinsights.com/industry-reports/artificial-intelligence-in-life-science-market#:~:North...

[19] https://www.coherentmarketinsights.com/industry-reports/artificial-intelligence-in-life-science-market#:~:In%20...

[20] https://intuitionlabs.ai/articles/generative-antibody-design-chai-3-pfizer#:~:In%20...

[21] https://www.nature.com/articles/s41586-026-10670-w#:~:We%20...

[22] https://www.nature.com/articles/s44386-026-00053-6#:~:Here%...

[23] https://www.nature.com/articles/s44386-026-00053-6#:~:This%...

[24] https://intuitionlabs.ai#:~:Throu...

[25] https://www.benchling.com/biotech-ai-report-2026#:~:This%...

[26] https://intuitionlabs.ai/articles/generative-antibody-design-chai-3-pfizer#:~:In%20...

[27] https://intuitionlabs.ai/articles/generative-antibody-design-chai-3-pfizer#:~:Chai-...

[28] https://www.benchling.com/biotech-ai-report-2026#:~:Data%...

[29] https://www.nature.com/articles/s41467-026-74501-2#:~:Emula...

[30] https://www.nature.com/articles/s41467-026-74501-2#:~:The%2...

[31] https://www.nature.com/articles/s41467-026-74501-2#:~:In%20...

[32] https://aiformanufacturing.org/case-studies/bristol-myers-squibb-cuts-npi-time-42-and-boosts-biologics-volume-40-with-ai-at-devens#:~:Deven...

[33] https://www.nature.com/articles/s41467-026-71204-6#:~:The%2...

[34] https://www.nature.com/articles/s41467-026-71204-6#:~:It%20...

[35] https://www.moduscreate.com/blog/scaling-ai-in-life-sciences#:~:Devia...

[36] https://www.moduscreate.com/blog/scaling-ai-in-life-sciences#:~:Four%...

[37] https://www.moduscreate.com/blog/scaling-ai-in-life-sciences#:~:Field...

[38] https://www.moduscreate.com/blog/scaling-ai-in-life-sciences#:~:Prior...

[39] https://www.zs.com/content/dam/global/master/pdfs/calculate-roi-for-ai-in-pharma-randd.pdf#:~:These...

[40] https://www.zs.com/content/dam/global/master/pdfs/calculate-roi-for-ai-in-pharma-randd.pdf#:~:Curre...

[41] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Among...

[42] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Meanw...

[43] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Other...

[44] https://www.benchling.com/biotech-ai-report-2026#:~:Devel...

[45] https://ardigen.com/ai-in-biotech-lessons-from-2025-and-the-trends-shaping-drug-discovery-in-2026/#:~:Acros...

[46] https://ardigen.com/ai-in-biotech-lessons-from-2025-and-the-trends-shaping-drug-discovery-in-2026/#:~:Integ...

[47] https://www.dlapiper.com/en-us/insights/publications/2025/01/fda-releases-draft-guidance-on-use-of-ai#:~:While...

[48] https://www.dlapiper.com/en-us/insights/publications/2025/01/fda-releases-draft-guidance-on-use-of-ai#:~:The%2...

[49] https://www.benchling.com/biotech-ai-report-2026#:~:30%25...

[50] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Achie...

[51] https://www.capgemini.com/wp-content/uploads/2026/01/120126CRI_Gen-AI-in-Lifesciences-_Final-interactive.pdf#:~:Our%2...

[52] https://www.capgemini.com/wp-content/uploads/2026/01/120126CRI_Gen-AI-in-Lifesciences-_Final-interactive.pdf#:~:Bioph...

[53] https://ispe.org/pharmaceutical-engineering/july-august-2024/artificial-intelligence-governance-gxp-environments#:~:AI%20...

[54] https://ispe.org/pharmaceutical-engineering/july-august-2024/artificial-intelligence-governance-gxp-environments#:~:Quali...

[55] https://ispe.org/pharmaceutical-engineering/july-august-2024/artificial-intelligence-governance-gxp-environments#:~:Estab...

[56] https://www.arxiv.org/pdf/2604.07512#:~:These...

[57] https://www.arxiv.org/pdf/2604.07512#:~:In%20...

[58] https://www.moduscreate.com/blog/scaling-ai-in-life-sciences#:~:Run%2...

[59] https://blogs.nvidia.com/blog/ai-in-healthcare-survey-2026/#:~:As%20...

[60] https://ispe.org/pharmaceutical-engineering/july-august-2024/artificial-intelligence-governance-gxp-environments#:~:Versi...

[61] https://www.biotech-asia.org/vol22no4/risk-based-validation-of-software-automation-and-artificial-intelligence-in-pharmaceuticals/#:~:Regul...

[62] https://www.biotech-asia.org/vol22no4/risk-based-validation-of-software-automation-and-artificial-intelligence-in-pharmaceuticals/#:~:Most%...

[63] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:In%20...

[64] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Accel...

[65] https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf#:~:Just%...

[66] https://blogs.nvidia.com/blog/ai-in-healthcare-survey-2026/#:~:70%25...

[67] https://blogs.nvidia.com/blog/ai-in-healthcare-survey-2026/#:~:For%2...

[68] https://blogs.nvidia.com/blog/ai-in-healthcare-survey-2026/#:~:%2A%2...

[69] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Notab...

[70] https://www.benchling.com/biotech-ai-report-2026#:~:%3E%2...

[71] https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf#:~:This%...

[72] https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf#:~:This%...

[73] https://ardigen.com/ai-in-biotech-lessons-from-2025-and-the-trends-shaping-drug-discovery-in-2026/#:~:This%...

[74] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Cyber...

[75] https://ispe.org/pharmaceutical-engineering/july-august-2024/artificial-intelligence-governance-gxp-environments#:~:Imple...

[76] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:The%2...

[77] https://www.precedenceresearch.com/artificial-intelligence-in-drug-discovery-market#:~:North...

[78] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Ninet...

[79] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Surve...

[80] https://www.capgemini.com/wp-content/uploads/2026/01/120126CRI_Gen-AI-in-Lifesciences-_Final-interactive.pdf#:~:In%20...

[81] https://www.prnewswire.com/news-releases/insilico-initiates-phase-iii-clinical-trial-for-rentosertib-its-ai-empowered-tnik-inhibitor-for-idiopathic-pulmonary-fibrosis-302819553.html#:~:CAMBR...

[82] https://www.prnewswire.com/news-releases/insilico-initiates-phase-iii-clinical-trial-for-rentosertib-its-ai-empowered-tnik-inhibitor-for-idiopathic-pulmonary-fibrosis-302819553.html#:~:Rento...

[83] https://www.bio-itworld.com/news/2026/07/08/insilico-medicine-launches-phase-iii-trial-for-ai-developed-drug#:~:%E2%8...

[84] https://www.bio-itworld.com/news/2026/07/08/insilico-medicine-launches-phase-iii-trial-for-ai-developed-drug#:~:It%20...

[85] https://www.prnewswire.com/news-releases/insilico-initiates-phase-iii-clinical-trial-for-rentosertib-its-ai-empowered-tnik-inhibitor-for-idiopathic-pulmonary-fibrosis-302819553.html#:~:Poten...

[86] https://aiformanufacturing.org/case-studies/bristol-myers-squibb-cuts-npi-time-42-and-boosts-biologics-volume-40-with-ai-at-devens#:~:Brist...

[87] https://aiformanufacturing.org/case-studies/bristol-myers-squibb-cuts-npi-time-42-and-boosts-biologics-volume-40-with-ai-at-devens#:~:The%2...

[88] https://aiformanufacturing.org/case-studies/bristol-myers-squibb-cuts-npi-time-42-and-boosts-biologics-volume-40-with-ai-at-devens#:~:Domai...

[89] https://aiformanufacturing.org/case-studies/bristol-myers-squibb-cuts-npi-time-42-and-boosts-biologics-volume-40-with-ai-at-devens#:~:Susta...

[90] https://www.fiercebiotech.com/biotech/takeda-inks-17b-ai-discovery-deal-iambic-therapeutics#:~:Tokyo...

[91] https://www.fiercebiotech.com/biotech/takeda-inks-17b-ai-discovery-deal-iambic-therapeutics#:~:%22As...

[92] https://www.fiercebiotech.com/biotech/takeda-inks-17b-ai-discovery-deal-iambic-therapeutics#:~:Iambi...

[93] https://www.microsoft.com/en/customers/story/26569-novo-nordisk-as-azure#:~:Befor...

[94] https://www.microsoft.com/en/customers/story/26569-novo-nordisk-as-azure#:~:%E2%8...

[95] https://www.microsoft.com/en/customers/story/26569-novo-nordisk-as-azure#:~:%22In...

[96] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:Promo...

[97] https://www.dlapiper.com/en-us/insights/publications/2025/01/fda-releases-draft-guidance-on-use-of-ai#:~:Furth...

[98] https://intuitionlabs.ai#:~:Exper...

[99] https://ardigen.com/ai-in-biotech-lessons-from-2025-and-the-trends-shaping-drug-discovery-in-2026/#:~:2025%...

[100] https://www.fiercebiotech.com/biotech/takeda-inks-17b-ai-discovery-deal-iambic-therapeutics#:~:Taked...

[101] https://www.fiercebiotech.com/biotech/takeda-inks-17b-ai-discovery-deal-iambic-therapeutics#:~:The%2...

[102] https://intuitionlabs.ai/articles/generative-antibody-design-chai-3-pfizer#:~:Pfize...

[103] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Achie...

[104] https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf#:~:Despi...

IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at https://intuitionlabs.ai/contact for a consultation.

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.