IntuitionLabs

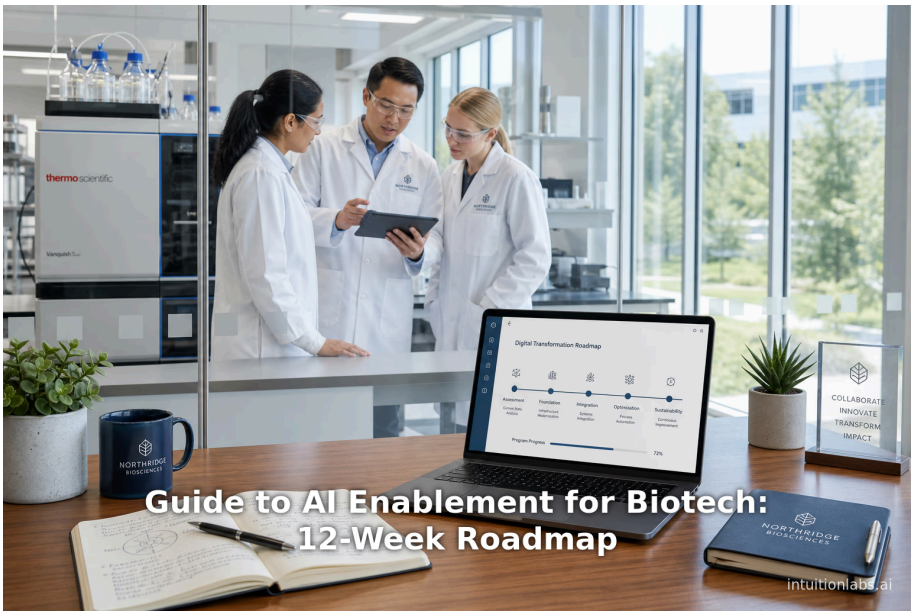
INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Guide to AI Enablement for Biotech: 12-Week Roadmap

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-06-28 · ai enablement · biotech ai · life sciences ai · ai drug discovery consulting · ai implementation roadmap · biotech ai consulting · artificial intelligence in life sciences · machine learning in biotech · biotech digital transformation · ai governance

Original article: https://intuitionlabs.ai/articles/ai-enablement-for-biotech-program



In this article

1. Executive Summary
2. Introduction and Background
3. AI Enablement Program Architecture: A 12-Week Roadmap
4. Use Cases and Functional Capabilities Across the R&D Value Chain
5. Market Share and Adoption Patterns
6. The AI Enablement and Consulting Provider Landscape
7. Data Analysis and Evidence
8. Case Studies and Real-World Examples
9. Novo Nordisk and NVIDIA: Sovereign AI Infrastructure for Drug Discovery
10. Eli Lilly and Insilico Medicine: AI-Native Biotech as R&D Partner
11. Illustrative Publisher-Reported Case: Workforce Enablement at a Clinical-Stage Biotech
12. Merck KGaA's AIDDISON Platform and Ginkgo's Virtual Cell Pharmacology Initiative
13. Implications and Future Directions
14. Frequently Asked Questions (FAQs)
15. Conclusion

Executive Summary

Biotech and biopharma organizations enter the second half of 2026 with a clear mandate to deploy artificial intelligence (AI) and a much less clear picture of how to do so without wasting money or breaking regulatory compliance. Industry surveys converge on the same tension: 82% of biopharma executives believe AI will fundamentally transform research and development (R&D) within five years, and a majority warn that companies failing to scale AI will fall behind in innovation and market relevance (capgemini.com). Yet the same research shows adoption is broad but shallow: only 45% of life sciences leaders report measurable AI improvement despite 71% saying deployment advanced in the past six months (deloitte.com), and a widely cited MIT NANDA analysis found that 95% of enterprise generative AI pilots fail to produce measurable profit-and-loss impact (fortune.com).

This report examines what a structured AI enablement program for biotech actually looks like: readiness assessment, workforce training, governance and policy, pilot selection, and measurement, sequenced across roughly a 12-week initial engagement with longer-term scaling. The market context is substantial. The global artificial intelligence in drug discovery market was valued at $2.3 billion in 2025 and is projected to grow from $2.9 billion in 2026 to $13.8 billion by 2033, a 24.8% compound annual growth rate (grandviewresearch.com). McKinsey estimates generative AI alone could generate $60 billion to $110 billion a year in economic value across pharmaceutical and medical-product industries (mckinsey.com). NVIDIA's 2026 survey of healthcare and life sciences organizations, whose respondents span digital healthcare, pharmaceutical and biotechnology, payers and providers, and medical technology, found 70% of respondents' organizations actively using AI, up from 63% in 2024, and 85% of executives saying AI is helping increase revenue (blogs.nvidia.com). For context, the U.S. Census Bureau's Business Trends and Outlook Survey found overall

AI usage across all U.S. businesses hovering between 17% and 20% over its December 2025 to May 2026 data (census.gov); because this covers a different population and measure than the NVIDIA sample, the two figures are not directly comparable.

The data also show why so many programs stall. Benchling's 2026 Biotech AI Report, drawn from a survey of roughly 100 biotech and pharmaceutical organizations, found that the single most commonly cited reason AI pilots fail is poor data quality, not model capability (benchling.com). MIT NANDA similarly found that purchasing and partnering for AI capability succeeds roughly 67% of the time, while fully internal builds succeed only about a third as often (fortune.com). Regulatory clarity is improving but still in motion: the U.S. Food and Drug Administration (FDA) issued its first draft guidance on AI in drug and biological product regulatory submissions in January 2025 (fda.gov), while the European Union's AI Act became applicable on 2 August 2026 with some exceptions, with high-risk systems embedded in regulated products granted an extended transition to 2 August 2028 (digital-strategy.ec.europa.eu).

Real-world deployments illustrate both the promise and the discipline required. Novo Nordisk's collaboration with NVIDIA and the Danish Center for AI Innovation gives its researchers access to the Gefion supercomputer for generative and agentic drug discovery workloads (nvidianews.nvidia.com), and Eli Lilly's expanded 2026 agreement with Insilico Medicine commits $115 million upfront plus up to $2.63 billion in milestones, for a total deal value of approximately $2.75 billion (insilico.com). These are far from isolated: Reuters counted more than 40 publicly disclosed AI-related pharma and biotech partnerships between January 2025 and May 2026, a dealmaking wave detailed further in the Data Analysis section below. At the workforce level, a life sciences AI consultancy such as IntuitionLabs runs enablement engagements built around a four-step assess, train, support, and measure framework, with workshop pricing that starts at $3,500 for a two-hour introductory session and $5,500 for a four-hour deep-dive format, detailed further below (these are the publisher's own published offering and prices, not independently verified).

For a biotech leadership team, the practical conclusion is that AI enablement is a change-management and data-governance program with a technology component, not the reverse. Organizations that succeed sequence readiness assessment before tool rollout, build tiered data classification before opening access to consumer AI tools, and measure adoption and time savings quarterly rather than declaring victory at go-live. Skipping this sequencing is likely to increase the risk of pilot failure, a risk reflected in the widely cited, though methodologically limited, 95% enterprise generative AI pilot failure statistic.

Introduction and Background

Biotech companies face a structural problem that predates generative AI: research and development productivity has been declining for decades. Eroom's Law, the informal but well-documented observation that the number of new drugs approved per inflation-adjusted dollar of R&D spending fell by roughly 50% every nine years between 1950 and 2010, describes an industry where each new therapy costs more to develop than the last. Widely cited industry estimates, which this report does not independently source, put drug development at a decade or more, costing billions of dollars per approved drug, with a high clinical-trial failure rate, in line with Grand View Research's characterization of traditional drug discovery timelines often exceeding a decade, with per-drug costs frequently surpassing $2 billion, a burden that AI-assisted

approaches such as drug repurposing aim to shorten and cheapen; the 3 to 12 year window and roughly $300 million average cited by Grand View for repurposing describe drug repurposing generally rather than AI-assisted programs specifically (grandviewresearch.com).

Against this backdrop, "AI enablement" has emerged as a distinct discipline from "AI research." Where AI research in biotech refers to building or licensing models for target identification, protein structure prediction, or molecule design, AI enablement refers to the organizational work of getting scientists, clinical operations staff, regulatory affairs teams, and commercial functions to actually use available AI tools safely, consistently, and in ways that generate measurable time savings. A biopharma organization can have access to the most capable foundation models in the world and still fail to capture value if its workforce lacks training, its data is siloed, and its governance policies either block adoption outright or, more commonly, fail to exist at all, pushing employees toward unsanctioned "shadow AI" use. Individual scientists have already moved ahead of their organizations on this front: surveyed R&D professionals across biotech now routinely reach for AI copilots and reasoning tools as a first stop when querying data, ahead of manual search methods, meaning the behavior change has already occurred among individual scientists even where formal enablement programs have not caught up. IQVIA frames the shift in similarly sweeping terms, with one of its machine learning vice presidents observing that "everything we do today as a pharma company or as a service provider to a pharma company, we can do better with AI across the lifecycle" (iqvia.com). Catalyst Group's own assessment of the market echoes this gap between individual enthusiasm and organizational readiness, noting that "most organizations in biotech, pharma, and AEC are at an early stage of AI adoption," aware of the potential but uncertain where to start (catalystgroup.ch).

The regulatory environment adds urgency and complexity in equal measure. The FDA's Center for Drug Evaluation and Research (CDER) reports a significant increase in the number of drug application submissions using AI components over the past several years, spanning nonclinical, clinical, postmarketing, and manufacturing phases (fda.gov), and it counted over 500 such submissions between 2016 and 2023 while developing its January 2025 draft guidance (fda.gov). That guidance, formally titled "Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products," proposes a seven-step, risk-based credibility assessment framework and remained open for public comment through April 7, 2025 (dlapiper.com). Notably, the guidance explicitly does not address AI models used in drug discovery itself, or AI used to streamline operations that do not affect patient safety, drug quality, or study reliability, meaning discovery-stage AI enablement work sits largely outside its scope for now (dlapiper.com). Where AI outputs do bear on regulatory decision-making, FDA's proposed framework treats models that make a final determination without human intervention as higher risk than those requiring human review before action, such as a model that flags out-of-specification manufacturing batches for human confirmation (dlapiper.com). CDER also established an internal AI Council in 2024 to coordinate policy, technology, and community-of-practice efforts across the agency (fda.gov), and the agency has continued refining its position, publishing Guiding Principles of Good AI Practice in Drug Development in January 2026 in coordination with the European Medicines Agency (fda.gov). CDER has also established formal external engagement pathways to work directly with sponsors and interested parties as AI use in submissions continues to expand (fda.gov). In the European Union, the AI Act entered into force on 1 August 2024 and became applicable on 2 August 2026, with some exceptions; the European Commission's AI Act page records that a political agreement reached on 7 May 2026 entered into force on 27 July 2026, with rules for high-risk AI systems embedded in

regulated products applying from 2 August 2028 and rules for systems used in certain other high-risk areas applying from 2 December 2027 (digital-strategy.ec.europa.eu). For a biotech operating on both sides of the Atlantic, this means an enablement program has to be designed against two moving regulatory targets simultaneously, not a single stable framework.

Pharmaceutical companies have long been in the vanguard of AI even before the recent surge of interest in generative tools, since researchers were already applying complex AI models to unlock the mechanisms of disease well before the current wave of enablement programs began (mckinsey.com). This report draws on primary survey data from Benchling, Capgemini, NVIDIA, Deloitte, McKinsey, and the U.S. Census Bureau; regulatory guidance and legal analysis from the FDA, the European Commission, DLA Piper, and Goodwin Procter; wire-service deal reporting from Reuters; and named deployment case studies from Novo Nordisk, Eli Lilly, and Insilico Medicine, alongside the service structures of active AI enablement and consulting providers serving the sector.

AI Enablement Program Architecture: A 12-Week Roadmap

A structured AI enablement program for a biotech organization generally moves through four phases that map to what one life sciences AI consultancy, IntuitionLabs, formalizes as assess, train, support, and measure, a sequence designed to move an organization "from AI-curious to AI-productive". While engagement lengths vary by scope, a practical initial roadmap can be compressed into roughly 12 weeks, with governance and measurement continuing indefinitely afterward. This sits between the shorter end of the market, exemplified by NEPF's 6 to 8 week function-level AI integration engagements (nepf.co), and the longer end, represented by multi-month strategic roadmaps such as the Launchpoint AI Roadmap co-developed by IntuitionLabs and Sakara Digital, which targets a 12 to 18 month phased plan anchored by a first 90-day action plan and inspired by Stanford's AI leadership model (sakaradigital.com). That roadmap is pitched specifically at life sciences teams who want a structured, compliant, value-first approach and who prefer a boutique, collaborative partner over a generic consultancy (sakaradigital.com), with deliverables that explicitly include governance, risk controls, and inspection-ready documentation alongside the use-case roadmap itself (sakaradigital.com).

Weeks 1 to 2: Readiness assessment. The program begins with an audit of current AI tool usage, data infrastructure, and organizational appetite for change. IntuitionLabs structures this as a free four-hour assessment covering current state analysis, review of the Microsoft ecosystem (licensing, Copilot eligibility, Teams and SharePoint configuration), team readiness evaluation, opportunity mapping, and a customized roadmap output. The output typically identifies the top 10 to 15 use cases where AI can save time or improve output quality, prioritized by impact, feasibility, and data sensitivity. The assessment is designed to be completed within roughly four hours and results in a customized roadmap at no cost to the organization, and the resulting plan typically distinguishes quick wins achievable within 30 days from medium-term goals at 90 days and a longer-term vision spanning 6 to 12 months. The assessment also maps departments by AI readiness level to help sequence which teams receive training first, since not every function is equally prepared to adopt new tools at the same pace. This phase matters disproportionately because data quality, not

model capability, is the leading cause of failed AI pilots in biotech R&D, as established above; a company that skips assessment and jumps straight to tool deployment is statistically more likely to end up in the enterprise-wide pilot failure cohort described in the Data Analysis section below.

Weeks 3 to 4: Governance and policy. Before broad tool rollout, regulated organizations need a data classification framework and usage policy. This step is not optional in practice: one industry survey, cited by IntuitionLabs (the publisher of this article) and not independently identified here, reportedly found that 68% of employees using AI at work have never disclosed it to their employer, which suggests informal, unmanaged AI use may already be occurring inside many biotech companies whether or not leadership has approved it. A representative framework uses four data tiers, Public, Confidential, Restricted, and Sensitive/Regulated, mapping each to specific AI tool permissions so that patient data, GxP records, regulatory submissions, and clinical trial data are explicitly barred from any AI tool while published research and press releases can be used freely. Developing a complete policy framework, including a master AI usage policy, data classification matrix, department-specific rules, and training materials, typically takes 4 to 6 weeks from kickoff to approved documents, though an interim baseline policy covering critical guardrails can be delivered in about 2 weeks. This framework must also be built to satisfy FDA and EU expectations directly: the FDA's draft guidance recommends that a credibility assessment plan and life cycle maintenance plan be tailored to model risk and "context of use," while the EU AI Act requires high-risk systems to maintain "adequate risk assessment and mitigation systems" before market placement (goodwinlaw.com) (digital-strategy.ec.europa.eu).

Weeks 5 to 8: Role-based training and pilot deployment. With guardrails in place, hands-on workshops train departments such as clinical operations, regulatory affairs, medical writing, data science, and commercial teams on department-specific use cases, prompt engineering, and quality verification, with the explicit goal of participants leaving with real, working AI projects rather than slide decks. This phase is also where organizations decide, use case by use case, whether to buy, build, or adapt, and most biotech organizations use a mix of purchased commercial applications, in-house builds, and adapted open-source tools rather than a single approach. MIT NANDA's research on this decision, cited earlier and detailed further in the Data Analysis section, is stark: purchased tools integrated through vendor partnerships succeed far more often than fully internal builds, largely because purpose-built vendor tools already encode workflow context that generic internal tools have to build from scratch.

Weeks 9 to 12: Support and early measurement. The final phase of the initial rollout establishes an ongoing support cadence, office hours, prompt library updates, and refinement of workflows, alongside the first measurement cycle: adoption rates, time saved, and early return on investment (ROI) signals. Deloitte's midyear 2026 life sciences outlook underscores why this step cannot be skipped: 71% of surveyed life sciences leaders said AI deployment had advanced at least somewhat in the prior six months, but only 45% reported measurable improvement from their AI initiatives, and just 13% reported measurable improvement at scale (deloitte.com). Deloitte's own recommendation for the remainder of 2026 is to make AI value measurable, focused on workflow redesign, governance, and key performance indicator definition rather than additional deployment for its own sake.

Beyond week 12, mature programs shift to quarterly reviews that reassess adoption, identify new opportunities, and expand AI usage across departments, since the enablement work does not end when the initial rollout does. This ongoing cadence matters because the AI tool landscape itself keeps shifting: Deloitte

found that 61% of surveyed life sciences leaders consider partnerships especially important specifically for building AI capabilities, more than for other categories of pipeline or product development partnership (deloitte.com).

Use Cases and Functional Capabilities Across the R&D Value Chain

AI enablement in biotech is not one initiative but a portfolio of use cases with very different maturity levels. Benchling's 2026 report, based on a November 2025 survey of approximately 100 biotechnology and pharmaceutical organizations already using AI in R&D, found the industry's "first killer apps" cluster around tasks with clean, easily validated data: literature and knowledge extraction (76% adoption), protein structure and property prediction (71%), scientific reporting and internal communication (66%), and target identification (58%). Adoption drops sharply in areas where data is scattered or hard to validate, including generative de novo drug design (42%), biomarker identification (40%), and absorption, distribution, metabolism, and excretion (ADME) modeling (29%) (benchling.com).

Capgemini's separate survey of 500 senior biopharma executives across nine countries corroborates this pattern from the discovery-phase perspective: target identification is the most widely adopted AI use case in drug discovery, used by 43% of organizations, and among that group, 32% report productivity improvements averaging a 28% time saving compared with pre-AI workflows (capgemini.com).

Beyond discovery, AI use cases extend across the full value chain:

- **Clinical trial optimization:** A majority of Capgemini's surveyed executives affirm generative AI can substantially improve clinical trial efficiency and outcomes, including endpoint prediction, dosing regimen selection, trial site selection, and adverse event prediction. IQVIA adds that AI in clinical development "increases the likelihood of success" by helping developers select the right indications, trial designs, and strategies for each asset (iqvia.com).
- **Regulatory submissions:** A strong majority of Capgemini's surveyed executives agree generative AI can fundamentally transform regulatory submission and approval workflows, with organizations already using AI for this purpose reporting productivity gains averaging close to a fifth in time saved.
- **Commercial and medical affairs:** Non-scientific use cases, such as document authoring, software engineering, and literature search, tend to see even higher adoption than scientific use cases, reflecting how much easier it is to validate outputs in these domains. IQVIA describes AI in commercialization as helping predict market shifts and deliver "the right message to the right stakeholder at the right time" (iqvia.com).
- **Agentic workflows:** IQVIA highlights agentic AI, systems that act autonomously on defined goals rather than simply generating outputs, as the next wave of enablement, citing a Gartner projection that by 2027, 50% of all business decisions will be augmented or automated by AI agents for decision intelligence (iqvia.com).

For biotech companies specifically, workforce-facing enablement often starts with general-purpose tools rather than specialized R&D platforms. IntuitionLabs' workshop curriculum, for instance, is built around ChatGPT Enterprise, Microsoft Copilot, Claude, Google Gemini, and specialized research tools like Deep Research, Perplexity, and Elicit, taught through department-specific modules covering clinical operations, regulatory affairs, medical writing, data science and biostatistics, quality assurance, and finance. This distinction, between enabling the existing workforce on general AI tools and licensing specialized scientific AI platforms, is one of the more consequential design decisions an enablement program makes, since the two paths require different budgets, timelines, and governance models.

Market Share and Adoption Patterns

Adoption of AI in biotech and life sciences has moved from experimentation to mainstream practice over the past two years, though the pace and depth of adoption vary considerably by function, company size, and geography. NVIDIA's second annual State of AI in Healthcare and Life Sciences survey, referenced earlier, found that 70% of respondent organizations are actively using AI and that 69% are using generative AI and large language models specifically, up from 54% the prior year (blogs.nvidia.com). Within the survey's pharmaceutical and biotechnology segment specifically, 57% report that drug discovery is being driven by AI, and 46% cite drug discovery and development as among their top return-on-investment use cases (blogs.nvidia.com). Open source adoption is also notably high, at 82% of respondents saying open source software and models are moderately to extremely important to their organization's AI strategy, alongside 47% of respondents using or assessing agentic AI (blogs.nvidia.com). These life sciences figures dwarf the general U.S. business population: the Census Bureau's biweekly survey found current AI use rates by firm size topping out around 37% even among the largest companies with 250 or more employees, well below the life sciences adoption levels reported by NVIDIA and Benchling (census.gov). The same Census Bureau data show the Information sector at 39.7% and Finance and Insurance at 33.9% AI use, both well above the 19.8% national rate but still far below life sciences-specific adoption figures (census.gov), and less than 20% of firms with four or fewer employees reported using AI at all (census.gov).

Benchling's data break adoption down further by respondent profile: its survey drew a majority of biopharma respondents alongside biotech respondents, weighted toward North America over Europe, and toward larger organizations over smaller ones. Notably, high-AI-adoption biotechs consistently report far greater wet lab and dry lab (computational) integration than low-adoption peers, discussed further below, suggesting that infrastructure integration and AI maturity reinforce each other.

Capgemini's executive-level view finds an even higher degree of strategic commitment: 79% of surveyed biopharma organizations are actively developing strategies to integrate AI across their R&D value chain, and 63% believe AI-driven platforms will originate most new molecular entities within the next decade, with the average share of AI-discovered targets in company pipelines projected to rise from 12% today to 34% in five years and 60% in ten (capgemini.com).

Deloitte's midyear 2026 outlook, based on an April 2026 survey of 150 life sciences executives combined with first-quarter earnings call analysis, adds a note of caution to these adoption figures: while 71% of leaders say AI deployment has advanced at least somewhat in six months and 35% report significant agentic AI progress including enterprise-wide rollout, only 45% report measurable performance improvement overall (deloitte.com).

An IDC survey of cross-industry AI adopters cited by IQVIA offers a complementary read: 73% of respondents reported "spectacular" or "significant" improvements in core operational processes from vendor applications embedded with AI, roughly in line with the optimism seen across life sciences-specific surveys (iqvia.com). IQVIA argues the real value of AI only emerges "when you connect intelligence across silos, clinical, regulatory and commercial, so insights flow seamlessly," a point directly relevant to biotech enablement programs still working through data fragmentation (iqvia.com).

The AI Enablement and Consulting Provider Landscape

Biotech leadership teams evaluating outside help for AI enablement encounter a fragmented but increasingly specialized market of consultancies, most of which position themselves specifically around regulated-industry compliance rather than generic AI training. *Table 1* below summarizes several active providers identified in this research, based on each firm's own public service descriptions.

Provider	Primary Focus	Engagement Model	Notable Emphasis
Neon Pi	AI governance and training for biotech and pharma	Strategy through execution, phased implementation	FDA 21 CFR Part 11, GxP compliance, AI governance frameworks, employee AI training (neonpi.com); also covers risk and compliance assessment as a discrete workstream (neonpi.com), alongside dedicated change management and employee AI training tracks (neonpi.com)
iuvo Technologies	IT-led AI consulting for biotech	Readiness assessment, strategic roadmapping, buildout	Combines IT expertise, scientific understanding, and a "white-glove" consulting approach for data workflows and lab automation (iuvotech.com); readiness assessment begins with "a comprehensive evaluation of your existing IT and data environments" (iuvotech.com), followed by strategic roadmapping that "defines your AI priorities, aligns with your compliance obligations" (iuvotech.com)
TeselaGen Catalyst	Software and AI consulting for biotech R&D	Full-stack build: data pipelines, ML/AI models, DevOps, lab systems	Targets the "AI implementation gap" between data science teams and production-grade MLOps in biotech (teselagen.com), citing disconnected knowledge silos across LIMS, ELNs, and lab instruments as a root barrier (teselagen.com), and a "bench-tool disconnect" where tools "don't fit into the actual workflows of your scientists" (teselagen.com)
Catalyst Group (Switzerland)	AI implementation for biotech, pharma, and engineering/construction	Readiness assessment through product development	Notes most biotech and pharma firms remain at an early AI adoption stage, uncertain where to start (catalystgroup.ch), builds predictive models for equipment failures, supply disruptions, and batch deviations (catalystgroup.ch), and deploys process automation "to reduce manual work, improve data quality, and accelerate decisions" (catalystgroup.ch)

Provider	Primary Focus	Engagement Model	Notable Emphasis
Sciagen AI	Life sciences technology and consulting partner	Advisory, process redesign, platform implementation	"Practitioners led, not consultant led," staffed by former clinical, regulatory, safety, and quality operators (sciagen.ai), spanning operating model design (sciagen.ai) and dedicated capability gap assessments (sciagen.ai)
NEPF	Function-level AI integration for life sciences leaders	6 to 8 week focused engagement	Explicitly positions against "enterprise-level" transformations priced in the millions, targeting single-function leaders instead (nepf.co), pitched at leaders who say any AI-driven change has to be defensible to leadership, quality, and potentially regulators (nepf.co)

An interpretive read of Table 1 shows the market splitting along two axes: general IT and software consultancies extending into biotech (iuvo, TeselaGen) versus life sciences domain specialists layering AI onto existing regulatory and clinical expertise (Neon Pi, Sciagen, NEPF, Catalyst Group). Pricing transparency is uneven across this market; most providers require a direct consultation to obtain a quote, in contrast to some workforce-training providers that publish workshop rates directly. For a biotech deciding between an IT-led build partner and a workforce enablement specialist, the underlying question is usually whether the near-term bottleneck is data engineering and MLOps capacity (favoring an IT-led build partner) or workforce adoption, governance, and change management (favoring an enablement specialist). Life sciences advisory firms with an adjacent AI enablement practice, such as IntuitionLabs, tend to frame their engagement specifically around the latter problem, emphasizing that generic AI training teaches prompting, while their own workshops are built around life sciences workflows with compliance guardrails from the first exercise, as described further in the workshop pricing table below.

Workshop-based enablement pricing, where published, follows a tiered structure. *Table 2* below summarizes one representative published pricing structure for comparison.

Format	Duration	List Price	Included Deliverables
Standard	2 hours	$3,500	Prompt engineering fundamentals, data classification overview, 2 to 3 hands-on exercises, role-specific prompt library, quick reference guide, session recording
Deep Dive (most popular)	4 hours	$5,500	Everything in Standard, plus advanced prompt engineering, custom GPT/Project building, Deep Research techniques, 5 to 6 exercises, personalized coaching, personal action plan
Train-the-Trainer	Multi-session	Custom quote	Everything in Deep Dive, plus facilitator training, complete facilitator guide, slide decks, practice delivery with feedback, ongoing facilitator support

This pricing structure suggests per-session workshop costs are a comparatively low-risk entry point relative to multi-month platform consulting engagements, letting a biotech test workforce enablement before committing to a larger data infrastructure or model-building program.

Data Analysis and Evidence

The quantitative case for structured AI enablement, rather than ad hoc tool adoption, rests on several converging data points from independent surveys and wire-service deal reporting conducted across 2025 and 2026.

Enterprise-wide pilot failure rates. The MIT NANDA initiative's "State of AI in Business 2025" report, based on 150 leadership interviews, a survey of 350 employees, and analysis of 300 public AI deployments, found that approximately 5% of AI pilot programs achieve rapid revenue acceleration while the large majority stall with little to no measurable profit-and-loss impact. More than half of surveyed enterprise generative AI budgets go to sales and marketing tools, even though the research found the largest realized ROI in back-office automation, a misallocation pattern with direct relevance to biotech, where the highest-value near-term use cases (literature review, scientific reporting, regulatory documentation) sit closer to back-office than to commercial functions (fortune.com).

Market sizing. Grand View Research values the global artificial intelligence in drug discovery market at $2.3 billion in 2025, projecting growth to $2.9 billion in 2026 and $13.8 billion by 2033, a 24.8% CAGR, with North America holding a 52.8% revenue share in 2025, drug optimization and repurposing leading with 52.46% of 2025 revenue by application, and pharmaceutical and biotechnology companies accounting for 59.19% of the market by end use, ahead of academic and research institutes (grandviewresearch.com). McKinsey's separate, broader estimate puts the annual economic value opportunity from generative AI across pharma and medical products at $60 billion to $110 billion, built from a bottom-up model of 63 individual generative AI use cases mapped across five industry domains, a figure informed by pre-generative-AI precedents such as AlphaFold2, ESMFold, and MoLeR using deep learning to predict the structures of nearly all known proteins (mckinsey.com).

The 2025 to 2026 AI dealmaking wave. Reuters compiled a running list of major AI-related pharma and biotech partnerships from January 2025 through May 2026 that illustrates how capital-intensive AI enablement has become at the largest organizations. AstraZeneca and CSPC Pharmaceutical structured a discovery deal worth $110 million upfront, up to $1.62 billion in development milestones, and up to $3.6 billion in sales milestones (reuters.com). Merck KGaA committed over $3 billion in potential value to Valo Health for target identification using human-data AI (reuters.com), Eli Lilly and NVIDIA committed up to $1 billion over five years to a closed-loop AI and robotics discovery lab, and Roche's acquisition of digital pathology AI firm PathAI carried $750 million upfront plus up to $300 million in additional milestones (reuters.com). Reuters also reported that Bristol Myers Squibb agreed to make Anthropic's Claude model available to more than 30,000 employees to accelerate drug discovery, development, and delivery, one of the clearest examples of enterprise-scale workforce enablement paired with a frontier AI vendor rather than a bespoke internal build (reuters.com).

Workforce and talent. Benchling's data show internal upskilling of existing scientific staff is the leading source of AI talent for biotech organizations, cited by 67% of respondents, well ahead of hiring from tech companies (21%) or management consulting firms (10%) (benchling.com). Independent of the biotech-specific data, a broader workplace survey cited by IntuitionLabs found that 68% of employees who use AI at work have never

disclosed that usage to their employer, a statistic with particular force in regulated industries where undisclosed AI use can create compliance exposure around protected health information, GxP records, and unreviewed AI-generated regulatory content.

Table 3 below consolidates the primary adoption and impact statistics referenced across these sources for quick comparison.

Metric	Value	Source
Organizations actively using AI (life sciences, 2026)	70%, up from 63% in 2024	NVIDIA survey (blogs.nvidia.com)
Executives believing AI will transform biopharma R&D	82%	Capgemini survey of 500 executives (capgemini.com)
Enterprise generative AI pilots delivering no measurable P&L impact	95%	MIT NANDA / Fortune (fortune.com)
Biotechs citing literature review/knowledge extraction as top adopted use case	76% adoption	Benchling 2026 Biotech AI Report (benchling.com)
AI drug discovery market size, 2026	$2.9 billion	Grand View Research (grandviewresearch.com)
Life sciences leaders reporting measurable AI improvement	45% (13% at scale)	Deloitte midyear 2026 outlook (deloitte.com)
Overall U.S. business AI usage (all sectors)	17% to 20%	U.S. Census Bureau BTOS (census.gov)

Read together, these figures describe an industry in the middle of a difficult but well-documented transition: strategic conviction is strong among leadership, tool usage is already widespread among individual scientists, and dealmaking activity is running into the tens of billions of dollars, but the organizational infrastructure, data governance, measurement discipline, and change management, needed to convert that usage into durable ROI lags well behind. This is precisely the gap that structured AI enablement programs are designed to close.

Case Studies and Real-World Examples

Novo Nordisk and NVIDIA: Sovereign AI Infrastructure for Drug Discovery

In June 2025, NVIDIA announced a collaboration with Novo Nordisk to accelerate drug discovery through the Danish Center for AI Innovation's (DCAI) Gefion supercomputer, built on NVIDIA DGX SuperPOD infrastructure (nvidianews.nvidia.com). Novo Nordisk uses NVIDIA BioNeMo for generative AI-powered drug discovery, NVIDIA NIM and NeMo microservices for agentic workflows, and NVIDIA Omniverse to build physically accurate simulation environments (nvidianews.nvidia.com). Research programs span using single-cell models to predict cellular responses to drug candidates, designing molecules with drug-like properties, and building biomedical large language models from Novo Nordisk's scientific literature (nvidianews.nvidia.com). Mishal Patel, senior vice president of AI and digital innovation at Novo Nordisk, described the goal as building

"custom models that will aid our scientists in developing new medicines faster and more efficiently," adding that "Gefion will allow us to run experiments at an unprecedented scale" (nvidianews.nvidia.com). This case illustrates enablement at the infrastructure and computational-science tier: large biopharma companies partnering with multiple hyperscale compute and model providers to build proprietary capability rather than relying on a single off-the-shelf tool, a materially different enablement path from workforce training programs but one that still depends on the same underlying data readiness and governance discipline.

Eli Lilly and Insilico Medicine: AI-Native Biotech as R&D Partner

Eli Lilly's relationship with Insilico Medicine, a clinical-stage biotechnology company built around generative AI and automation, began as a research collaboration announced in November 2025 (insilico.com) and expanded substantially in March 2026 into a global R&D and licensing agreement (insilico.com). Under the March 2026 agreement, Lilly received an exclusive worldwide license for development, manufacturing, and commercialization of novel oral therapeutics discovered using Insilico's Pharma.AI platform, with Insilico eligible for a $115 million upfront payment plus development, regulatory, and commercial milestones bringing the total potential deal value to approximately $2.75 billion, in addition to tiered royalties on future sales (insilico.com). Insilico founder and CEO Alex Zhavoronkov framed the partnership around scale of target discovery, stating that the company can "identify multi-purpose targets driving multiple diseases at the same time" by deploying frontier AI technologies across biomarkers and world models of human and animal biology (insilico.com). Taken together, this case represents the opposite end of the enablement spectrum from Novo Nordisk's infrastructure buildout: rather than building internal AI capability from scratch, Lilly is buying access to an AI-native biotech's discovery engine, echoing the MIT NANDA finding described earlier that partnership-based AI adoption outperforms internal builds by a wide margin.

Illustrative Publisher-Reported Case: Workforce Enablement at a Clinical-Stage Biotech

Illustrating the workforce training end of AI enablement rather than the platform-building end, a clinical-stage biotech company (identified by its consultancy only by stage and sector, not by name) engaged IntuitionLabs for two AI workshops covering its cross-functional team. Within weeks of the engagement, team members were independently building custom ChatGPT Projects for their own daily workflows and running Deep Research sessions for competitive intelligence and regulatory landscape analysis, activities they had not previously undertaken with AI tools. This case is representative of the workshop-driven enablement model described earlier in this report: two structured sessions, built around role-specific exercises rather than generic prompting instruction, were reported by IntuitionLabs, this article's publisher, to have changed working habits within weeks; this is a single, anonymous, self-reported example and not independent evidence of durable behavior change.

Merck KGaA's AIDDISON Platform and Ginkgo's Virtual Cell Pharmacology Initiative

Beyond single-company partnerships, the broader industry has invested in proprietary AI tooling for drug discovery. On December 5, 2023, MilliporeSigma, the U.S. and Canada life science business of Merck KGaA, Darmstadt, Germany, launched its AIDDISON drug discovery software, which the company itself described as the first software-as-a-service platform that bridges the gap between virtual molecule design and real-world manufacturability through Synthia retrosynthesis software application programming interface (API) integration, combining generative AI, machine learning, and computer-aided drug design (emdgroup.com). More recently, in November 2025, Ginkgo Datapoints launched the Virtual Cell Pharmacology Initiative, an open-source platform for standardized virtual cell modeling that aims to test over 100,000 compounds and generate more than 12 billion data points using its V-Ref293 engineered cell line and DRUG-seq RNA profiling method (grandviewresearch.com). Both cases underscore that even the largest, best-resourced organizations are still building foundational AI-native infrastructure rather than treating deployment as complete, reinforcing the report's broader argument that AI enablement is an ongoing capability-building program, not a one-time technology purchase.

Implications and Future Directions

Several forward-looking implications follow from the evidence gathered in this report. First, regulatory timelines will increasingly constrain how biotech organizations sequence AI rollout across jurisdictions. With the EU AI Act's transparency rules taking effect in August 2026 and high-risk system obligations phasing in through 2 December 2027 and 2 August 2028 depending on category (digital-strategy.ec.europa.eu), multinational biotechs will need governance frameworks that satisfy the stricter of two regimes rather than maintaining parallel compliance tracks. On the U.S. side, the FDA's draft guidance includes a hypothetical scenario in which an AI model categorizes patients by adverse-event risk to determine inpatient versus outpatient monitoring, illustrating why the agency treats any AI model making a final determination without human intervention as inherently higher risk (goodwinlaw.com). Legal analysts covering the guidance note that it would be the first FDA guidance of its kind to incorporate comprehensive recommendations for the design, development, documentation, and maintenance of AI models used to support regulatory decisions, if finalized (goodwinlaw.com). The same analysis notes the guidance arrived just weeks after a change in U.S. presidential administration explicitly aimed at removing barriers to American AI leadership, underscoring how quickly the domestic policy backdrop for AI in drug development has shifted within a single year (goodwinlaw.com).

Second, the shift toward agentic AI, systems capable of autonomous, multi-step action rather than single-turn generation, will raise the stakes on governance frameworks that were originally designed for simpler chatbot-style tools. IQVIA's citation of a Gartner projection that half of all business decisions will be AI-agent-augmented by 2027 implies that the four-tier data classification frameworks common in current biotech AI policies will need extension to cover autonomous action, not just content generation, within the next one to two years (iqvia.com). IQVIA also notes an important adoption asymmetry: nonexpert user groups tend to adopt agentic AI tools faster than domain experts, since experts have a lower tolerance for the tools' current

limitations, a dynamic biotech enablement programs will need to manage explicitly rather than assume expertise correlates with faster adoption (iqvia.com). IQVIA frames the underlying shift as agentic AI functioning as "automation at scale, freeing experts to focus on strategy," and cautions that successful adoption still depends on human review for consequential or sensitive decisions even as automation expands (iqvia.com).

Third, community sentiment among practicing biotech scientists remains more skeptical than executive survey data suggests, a gap enablement programs should not ignore. In discussion threads on Reddit's r/biotech community, practitioners have pushed back on aggressive AI timelines, with one commenter distinguishing sharply between "the AI used in drug discovery and the AI you ask to write your email," arguing that confusion between the two inflates expectations for near-term scientific breakthroughs (reddit.com). Another commenter in the same thread argued that "discovery is not the problem, clinical trials are the sink both in terms of time and money," pointing to failed oncology trials as a bigger cost driver than any discovery-stage inefficiency AI has addressed to date (reddit.com). A third participant offered a more measured middle position, arguing that AI's real contribution is putting "structure to unstructured or messy information," which is itself much of what drug discovery involves, even if it is not "coming up with a new drug by walking up to it and asking for a drug for Alzheimer's" (reddit.com). These figures and framings come from unverified forum commentary rather than a research firm and should be read as illustrative of practitioner sentiment, not as validated statistics; they nonetheless signal where scientific staff believe enablement programs should focus, closer to trial design and execution than to discovery-stage tooling alone (reddit.com). One commenter summarized frontline frustration bluntly, describing repeated leadership town halls "basically been begging for people to come up with ways to use AI" only to land on low-value examples, a pattern enablement programs are specifically designed to prevent by anchoring adoption in real departmental workflows rather than generic mandates (reddit.com).

Fourth, talent strategy will remain a binding constraint. Because internal upskilling remains the dominant source of biotech AI talent, as noted earlier, enablement programs that under-invest in structured internal training in favor of one-time external hires are likely to under-deliver relative to organizations treating upskilling as a continuous program. Finally, the measurement gap identified earlier in this report, where AI deployment consistently outpaces the ability to demonstrate value, is likely to become the primary differentiator between biotech organizations that sustain AI investment through the current budget cycle and those that see funding pulled back once initial enthusiasm fades.

Frequently Asked Questions (FAQs)

What is AI enablement for biotech, and how is it different from AI research? AI enablement refers to the organizational work of training staff, building governance policy, classifying data, and measuring adoption so that available AI tools are used safely and effectively across a biotech organization, as distinct from AI research, which focuses on building or licensing the underlying models themselves.

How long does a biotech AI implementation roadmap take? Initial enablement engagements range from about 6 to 8 weeks for a single function (nepf.co) to a compressed 12-week program covering assessment, governance, training, and initial measurement, with strategic roadmaps extending 12 to 18 months when they include a first 90-day action plan for company-wide rollout (sakaradigital.com).

What does an AI strategy for biotech companies need to address that a general enterprise AI strategy does not? It must address FDA and EU AI Act compliance requirements, data classification for protected health information and GxP records, validation and audit trail requirements for AI-generated content used in regulated submissions, and department-specific workflows in clinical operations, regulatory affairs, and medical writing that general AI training does not cover.

Where is machine learning already delivering measurable results in biotech? The strongest current results, as detailed earlier in this report, are in literature review and knowledge extraction, protein structure and property prediction, and scientific reporting, all areas where data is relatively clean and outputs are easy to validate.

What does life sciences AI enablement cost? Published workshop-based training pricing, detailed in Table 2 above, starts around $3,500 for a two-hour introductory session and $5,500 for a four-hour deep-dive format, with custom pricing for multi-session train-the-trainer engagements. These are the publisher's own published prices and have not been independently verified. Larger platform-building or AI research collaborations, such as pharma-biotech R&D partnerships, can run into hundreds of millions or billions of dollars in upfront and milestone payments, as seen in deals like AstraZeneca's agreement with CSPC Pharmaceutical and Merck KGaA's agreement with Valo Health (reuters.com).

Should a biotech company hire an AI drug discovery consulting firm, or build internal capability? The evidence favors a hybrid approach weighted toward partnership over pure internal build: MIT NANDA's research found purchased or partnership-based AI adoption succeeds roughly 67% of the time compared with about a third as often for fully internal builds (fortune.com), while internal upskilling, as discussed in the Data Analysis section, remains the leading source of AI talent even at organizations that do partner externally.

How does biotech digital transformation differ from AI enablement specifically? Digital transformation is the broader modernization of data infrastructure, wet-dry lab integration, and IT systems, while AI enablement is a narrower, faster-moving workstream layered on top of that infrastructure, focused specifically on workforce training, governance, and tool adoption. Benchling's data show that biotechs with more mature digital integration, measured by wet-dry lab connectivity, also show meaningfully higher AI adoption: 30% of high-AI-adoption biotechs describe their wet-dry lab workflows as highly integrated, compared with only 18% among low-adoption organizations (benchling.com).

Is AI drug discovery consulting regulated, and does the FDA require sign-off before use? The FDA does not pre-approve AI tools generically, but its draft guidance recommends early engagement so sponsors and the agency can agree on the appropriate credibility assessment activities for a given model risk and context of use before a regulatory submission relies on AI-generated data (dlapiper.com), a recommendation legal analysts note is intended to help set expectations regarding appropriate credibility assessment activities before problems surface late in development (dlapiper.com).

Conclusion

The evidence assembled in this report points to a consistent conclusion: AI enablement for biotech succeeds or fails based on organizational discipline, not model capability. Strong executive conviction that AI will transform biopharma R&D coexists with a widely cited finding that 95% of enterprise generative AI pilots

(across industries, based on a limited-methodology MIT NANDA study) fail to show measurable profit-and-loss impact; the evidence does not demonstrate what drives that gap, but surveys point to data quality, governance, and measurement as plausible contributors, which is why this report recommends investing in them before deploying tools. A structured roadmap, readiness assessment, tiered data governance, role-based training, and quarterly measurement, compressed into an initial 12-week engagement and sustained indefinitely afterward, addresses the specific failure modes that independent research from Benchling, Capgemini, Deloitte, and MIT NANDA all converge on: poor data quality, undisclosed shadow AI use, misallocated budgets, and an absence of success metrics.

The market context makes the stakes concrete. A drug discovery AI market growing from $2.9 billion in 2026 toward $13.8 billion by 2033, deal structures like Lilly's $2.75 billion collaboration with Insilico Medicine, and more than 40 additional disclosed AI partnerships tracked by Reuters across pharma, diagnostics, and biotech all demonstrate that capital is already flowing at scale into AI-native biotech capability. For the majority of biotech organizations that will never build a Gefion-scale supercomputing partnership or sign a billion-dollar AI licensing deal, the more immediately actionable path is workforce enablement: readiness assessment, governance, department-specific training, and disciplined measurement, delivered either internally or through a specialist consultancy, sequenced to avoid the pilot-failure patterns the data describe so clearly. Regulatory frameworks on both sides of the Atlantic are still being finalized, but their direction is set: AI use in regulated drug development will require documentation, validation, and human oversight regardless of which specific deadline applies in a given jurisdiction. Biotech organizations that build that discipline into their AI enablement program now, rather than retrofitting it after a compliance incident or a failed pilot, are the ones positioned to convert today's strong AI conviction into the kind of measurable, sustained value that remains, as of mid-2026, still the exception rather than the rule.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.capgemini.com/wp-content/uploads/2026/01/120126CRI_Gen-AI-in-Lifesciences-_Final-interactive.pdf
2. https://www.deloitte.com/us/en/insights/industry/health-care/midyear-2026-life-sciences-outlook.html
3. https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo
4. https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-drug-discovery-market
5. https://www.mckinsey.com/industries/life-sciences/our-insights/generative-ai-in-the-pharmaceutical-industry-moving-from-hype-to-reality
6. https://blogs.nvidia.com/blog/ai-in-healthcare-survey-2026
7. https://www.census.gov/library/stories/2026/05/ai-use-businesses.html
8. https://www.benchling.com/biotech-ai-report-2026
9. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological
10. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
11. http://nvidianews.nvidia.com/news/nvidia-partners-with-novo-nordisk-and-dcai-to-advance-drug-discovery
12. https://insilico.com/news/uiy12zcjg1-insilico-medicine-announces-global-rampd

13. https://www.iqvia.com/blogs/2025/10/how-emerging-ai-capabilities-are-reshaping-life-sciences
14. https://catalystgroup.ch/service-ai
15. https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development
16. https://www.dlapiper.com/en-us/insights/publications/2025/01/fda-releases-draft-guidance-on-use-of-ai
17. https://www.nepf.co/ai-lifesciences
18. https://sakaradigital.com/launchpoint-ai-roadmap
19. https://www.goodwinlaw.com/en/insights/publications/2025/01/alerts-lifesciences-aiml-fda-publishes-its-first-draft-guidance
20. https://www.iqvia.com/blogs/2026/01/exploring-the-value-of-agentic-ai-in-life-sciences
21. https://neonpi.com
22. https://www.iuvotech.com/ai-consulting/biotech
23. https://teselagen.com/catalyst
24. https://www.sciagen.ai/life-sciences
25. https://www.reuters.com/business/healthcare-pharmaceuticals/pharma-sector-doubles-down-ai-amid-hopes-slashing-costs-timelines-2026-05-20
26. https://insilico.com/news/e9xm9y0if1-insilico-and-lilly-enter-a-research-and
27. https://www.emdgroup.com/en/news/aiddison-drug-discovery-software-05-12-2023.html
28. https://www.reddit.com/r/biotech/comments/1u77rwj/is_ai_genuinely_transforming_biotech_industry_or

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.