IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Biotech AI Strategy: How to Prioritize Initiatives in 2026

> 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 · 2026-06-23 · biotech ai strategy · ai in life sciences strategy · ai prioritization framework · pharma ai investment · drug discovery ai roi · ai roadmap biotech · life sciences ai consulting · fda ai guidance · biotech ai use cases

Original article: https://intuitionlabs.ai/articles/biotech-ai-strategy-prioritization-framework



In this article

1. Executive Summary
2. Introduction and Background
3. Building the Prioritization Matrix: The Five Scoring Axes
4. The AI Maturity Framework for Drug Discovery Use Cases
5. Regulatory Risk Tiering: The FDA and EMA Frameworks
6. Data Readiness as the Binding Constraint
7. Sequencing: Fast-Track Wins Versus Long-Cycle Bets
8. Data Analysis and Evidence
9. Case Studies and Real-World Examples
10. Insilico Medicine: Generative Chemistry to Global Pharma Partnership
11. Recursion Pharmaceuticals and Exscientia: Consolidation as Strategy
12. Xaira Therapeutics: A Billion-Dollar AI-Native Launch
13. Sanofi: AI-Governed Portfolio Decision-Making at Enterprise Scale
14. Pfizer: Embedding AI Across an Existing Commercial Pipeline
15. Implications and Future Directions
16. Frequently Asked Questions (FAQs)
17. Conclusion
18. References

Executive Summary

Biotech and pharmaceutical executives no longer lack candidate artificial intelligence (AI) projects; they lack a defensible way to choose among dozens of competing ones. As of July 2026, industry survey data converge on a consistent picture: enthusiasm for AI is nearly universal, but conversion of that enthusiasm into measurable value remains rare. Deloitte's 2026 Life Sciences Outlook Survey found that only **22%** of life sciences leaders say they have successfully scaled AI, and just **9%** report achieving significant returns on their investments (www.deloitte.com). Separately, ZS's 2026 CDIO survey of 115 pharma and biotech technology executives found that just **40% of AI pilots** that clear an initial value bar make it to scaled deployment, and **68%** of respondents identified neglecting data quality and governance early as the main reason AI initiatives fail (www.zs.com).

This report lays out a practical framework for prioritizing biotech AI strategy: how organizations should select, sequence, and govern AI initiatives across research and development (R&D), clinical development, manufacturing, and commercial operations. The core analytical tool is a multi-axis scoring approach, exemplified by consulting frameworks that score candidate use cases on **strategic value, technical feasibility, data readiness, regulatory risk, and time-to-value**, typically weighted at approximately 25%, 20%, 20%, 15%, and 20% respectively (detailed later in this report). A parallel framework used in drug-discovery-specific evaluation scores AI maturity across five criteria: the right problem, the right data, the right

capabilities, the right market configuration, and demonstrated impact (pharmaphorum.com). Both converge on the same conclusion: data readiness, not algorithmic sophistication, is usually the binding constraint on realized value.

Quantitatively grounded use cases already deliver: Benchling's 2026 Biotech AI Report, based on a November 2025 survey of roughly 100 biotechnology and pharmaceutical organizations, found literature and knowledge extraction reaches **76% adoption**, protein structure prediction **71%**, and target identification **58%**, with **50%** of respondents reporting faster time-to-target today (www.benchling.com). AI-guided screening delivers hit rates of **22 to 46%** against a roughly **2%** baseline for random screening (publisher's own synthesis, not an independent primary source), and a peer-reviewed analysis in *Drug Discovery Today* found AI-discovered molecules achieved an **80 to 90%** Phase I success rate, well above historic industry averages, though the Phase II rate of roughly **40%** was comparable to industry norms (pubmed.ncbi.nlm.nih.gov).

The regulatory landscape has matured substantially. The U.S. Food and Drug Administration (FDA) issued draft guidance in January 2025 establishing a seven-step, risk-based credibility framework for AI models used in regulatory submissions (www.fda.gov), and in January 2026 the FDA and the European Medicines Agency (EMA) jointly published ten Guiding Principles of Good AI Practice in Drug Development that scale oversight proportionally to risk (detailed later in this report). Capital continues to flow toward AI-native biotech: Xaira Therapeutics launched in April 2024 with **$1 billion** in committed funding (detailed in the case studies below), Recursion Pharmaceuticals acquired Exscientia in an all-stock deal valued at **$688 million** in August 2024 (www.reuters.com), and Eli Lilly signed a collaboration with Insilico Medicine worth up to **$2.75 billion** including a $115 million upfront payment, announced in 2026 (insilico.com). The global AI-in-drug-discovery market is projected to grow from approximately **$7.62 billion in 2026** to **$17.81 billion by 2035**, a compound annual growth rate (CAGR) of **9.90%**, according to Precedence Research (www.precedenceresearch.com).

The report also examines organizational and governance factors: how firms should build cross-functional prioritization committees, structure data foundations, sequence commercial versus scientific investments, and evaluate vendor and consulting partners, including platform providers such as Veeva Systems. The analysis concludes that successful biotech AI strategy treats prioritization not as a one-time exercise but as a recurring governance discipline, anchored in honest data-readiness assessment, risk-tiered oversight aligned to the FDA/EMA principles, and portfolio management practices that explicitly track the gap between piloted and scaled initiatives, a discipline that remains rare given that **63%** of data management leaders surveyed by Gartner either lack, or are unsure whether they have, the right data practices to support AI (www.gartner.com).

Introduction and Background

Artificial intelligence has moved from experimental curiosity to board-level agenda item across the biotechnology and pharmaceutical sectors. Nearly half of surveyed life sciences executives, **48%**, identified accelerated digital transformation as a trend likely to have a substantial impact on their organizations in 2026, a statistically significant increase compared with 2025, according to Deloitte's 2026 Life Sciences Outlook survey (www.deloitte.com). Generative AI specifically was cited by **41%** of respondents as an influential trend, and **30%** cited agentic AI, defined as AI systems that can act autonomously to achieve goals, make decisions, and perform tasks, a category new to the survey this year.

Yet enthusiasm has not translated evenly into results. ZS's October 2025 CDIO outlook survey of 115 pharmaceutical and biotechnology technology executives found that nine in ten (approximately **90%**) see competitive and regulatory pressures, including AI disruption, as active threats to business growth. Top priorities to sustain growth include accelerated discovery, cited by **52%** of respondents, and patient engagement, cited by **43%** (www.zs.com). This intensity of expectation, combined with limited realized returns, is precisely the environment in which prioritization discipline becomes a binding organizational constraint rather than a nicety.

Confidence varies sharply by geography and function, which itself has implications for how centrally an AI portfolio should be governed. Deloitte's 2026 Life Sciences Outlook Survey found that more than **75%** of responding biopharma and medtech executives are confident in their own organizations' financial outlooks for the coming year, yet only **41%** feel optimistic about the health of the global economy overall (www.deloitte.com). Regulatory change is itself now a named strategic trend, not a background condition: one in two non-US respondents, **51%**, pointed to national regulatory changes such as the European Union's AI Act as a factor likely to affect market access, pricing, and reimbursement in 2026, and 35% of respondents flagged rising cybersecurity concerns as a factor shaping strategy.

The problem is structural. Most biotech and pharma organizations already have more candidate AI use cases than they have budget, talent, or data infrastructure to support. A boutique digital strategy consultancy, Sakara Digital, frames the challenge directly: "Most pharma and biotech leaders working through their AI portfolio do not lack candidate use cases. They lack a defensible mechanism for choosing among the candidates" (sakaradigital.com). Gartner predicted in July 2024 that at least **30%** of generative AI projects would be abandoned after proof of concept by the end of 2025 due to poor data quality, inadequate risk controls, escalating costs, or unclear business value, which underscores the stakes of getting prioritization wrong (www.gartner.com). Gartner's own newsroom frames the data-readiness dimension of that risk even more starkly: in a February 2025 release based on a July 2024 survey of 1,203 data management leaders, the firm states that "through 2026, organizations will abandon **60%** of AI projects unsupported by AI-ready data," and that **63%** of surveyed organizations either lack, or are unsure whether they have, the right data management practices for AI (www.gartner.com). The two figures, 30% and 60%, are not directly comparable since they measure different scopes, generative AI projects abandoned after proof of concept against a broader set of failure causes, versus AI projects specifically unsupported by AI-ready data, but both point in the same direction: data readiness failures, not model quality, are the dominant risk to a biotech AI portfolio.

This report addresses the question of biotech AI strategy from a practitioner's perspective: how should a chief digital officer, head of R&D informatics, or portfolio management leader decide which AI initiatives to fund first, how should those decisions be structured to survive board-level scrutiny, and what organizational, data, and regulatory prerequisites must be in place before an initiative can responsibly scale. The analysis draws on primary survey data from Deloitte, ZS Associates, and Benchling; regulatory guidance from the FDA and EMA; peer-reviewed clinical outcome data; and named case studies spanning Sanofi, Pfizer, Insilico Medicine, Recursion Pharmaceuticals, Exscientia, and Xaira Therapeutics. Because IntuitionLabs works with life sciences organizations as a **Veeva X-Pages Partner** focused on Vault CRM and adjacent commercial and compliance systems, this report also situates AI prioritization within the practical reality of platform ecosystems that most biotech and pharma commercial and quality organizations already depend on.

The remainder of this report proceeds as follows. The next section walks through the five-axis prioritization matrix that consulting practices use with pharma and biotech clients, followed by a section on the parallel AI-maturity framework used specifically for drug discovery use case evaluation. Subsequent sections cover regulatory risk tiering under the FDA and EMA frameworks, data readiness as the binding operational constraint, and the organizational and governance model required to move initiatives from pilot to scaled production. A dedicated data section aggregates the quantitative evidence on adoption, ROI, and market sizing. Case studies examine five named organizations. The report closes with implications for the next 12 to 24 months, frequently asked questions, and a conclusion.

Building the Prioritization Matrix: The Five Scoring Axes

The single most consequential decision in biotech AI strategy is not which technology to adopt, but which candidate initiative to fund first. A structured scoring matrix disciplines that judgment by making the criteria and trade-offs explicit rather than implicit. The most widely referenced version of this approach, used by boutique life sciences digital consultancies with pharma and biotech clients, scores each candidate use case against five axes (sakaradigital.com):

- **Strategic value**: the contribution to enterprise strategic priorities if the use case succeeds, scored not on the size of the addressable benefit but on alignment with what leadership has already committed to.
- **Technical feasibility**: the likelihood the use case can be delivered with available capabilities, in-house or vendor, within a reasonable timeline; this is the axis where AI advocates most often overscore, so healthy skepticism is warranted.
- **Data readiness**: the availability, quality, governance, and accessibility of the data the use case requires. Data readiness is consistently identified as the most frequently cited blocker to enterprise AI value capture, ahead of both algorithmic sophistication and talent gaps.
- **Regulatory risk**: the regulatory exposure the use case creates, scored on an inverted scale where lower risk scores higher. A use case affecting safety, quality, or efficacy carries materially different regulatory weight than one confined to marketing or human resources.
- **Time-to-value**: the expected payback period and the magnitude of value at the end, combining speed and scale because either alone is an insufficient basis for prioritization.

Each axis is scored on a **1 to 5 scale** with anchored definitions per score level; without anchoring, scores drift toward the middle and the matrix loses its discriminative power. *Table 1* below summarizes representative anchor definitions used in practice.

Table 1. Representative Scoring Anchors for the Five-Axis Prioritization Matrix

Axis	Score 1 (Low)	Score 3 (Moderate)	Score 5 (High)
Strategic value	Tactical efficiency; not on leadership's stated priority list	Supports a stated strategic priority but is not the primary mechanism	Directly advances a top-three strategic priority

Axis	Score 1 (Low)	Score 3 (Moderate)	Score 5 (High)
Technical feasibility	Requires novel research or unproven architecture	Uses proven architecture but requires meaningful customization	Uses well-understood architecture with available capability
Data readiness	Required data does not exist or requires major collection effort	Required data exists but needs significant preparation	Required data is available, governed, and accessible
Regulatory risk (inverted)	High-impact AI directly affecting safety, quality, or efficacy	Moderate exposure with an established validation pathway	Low or no regulatory exposure
Time-to-value	Multi-year payback or unclear value mechanism	12 to 18 months to material value	Less than 12 months to material value at meaningful scale

(Source: sakaradigital.com)

The table above collapses five distinct evaluation dimensions into a single, comparable rubric, which is the point: it forces a technical team's confidence in feasibility to be weighed openly against a regulatory affairs team's assessment of risk, rather than letting each function's priorities silently dominate the discussion. This structure is especially important because default axis weights, typically **25% strategic value, 20% technical feasibility, 20% data readiness, 15% regulatory risk, and 20% time-to-value**, are adjusted based on organizational context (sakaradigital.com). Clients under board-level revenue pressure often raise the time-to-value weight to 25 to 30%; clients in early regulatory engagement cycles often raise the regulatory risk weight to 20 to 25%; and clients with significant data infrastructure gaps often raise the data readiness weight to 25%, reflecting that data work is the binding constraint for their portfolio.

The weighted sum produces a score on a 1 to 5 scale, which practitioners then convert into tiers (A: 4.0 and above; B: 3.0 to 3.9; C: 2.0 to 2.9; D: below 2.0) rather than reporting the raw number, because precise decimal scores invite false confidence while tiers create natural cut lines for portfolio decisions. A calibration session, run with the AI governance committee or a designated cross-functional group before any use case is scored, walks through three to four example use cases collectively to surface implicit assumptions and correct the two most common biases: technical advocates systematically overscoring feasibility, and business advocates systematically overscoring strategic value.

The AI Maturity Framework for Drug Discovery Use Cases

A parallel framework, specific to evaluating AI applications inside drug discovery rather than the broader enterprise portfolio, measures maturity across five criteria: **the right problem**, which determines the suitability and advantages of AI over traditional approaches; **the right data**, which evaluates current data availability, accessibility, and quality; **the right capabilities**, which examines available tools, platforms, and expertise; **the right market configuration**, which considers the number of active players and the extent of collaboration interest from pharmaceutical companies; and **demonstrated impact**, which evaluates AI's ability to generate novel insights and identifies tangible successes achieved thus far (pharmaphorum.com).

This framework is instructive when applied to specific use case families. In small-molecule design, companies such as Exscientia (now part of Recursion) and Insilico Medicine are cited as leading efforts, with AI-designed small molecules progressing into Phase 2 clinical trials, expected to provide the field's clearest evidence yet of AI maturity in drug design (pharmaphorum.com). Antibody design, by contrast, remains earlier stage: despite Xaira Therapeutics securing over **$1 billion** in funding specifically to focus on de novo antibody design, the article notes that difficulties persist, even as growing pharma investment and disclosed partnerships are expected to drive maturity forward (pharmaphorum.com).

The practical implication for strategy teams is that "the right problem" and "demonstrated impact" criteria should be evaluated separately by therapeutic modality (small molecule, antibody, cell and gene therapy) rather than treated as a single undifferentiated "AI in drug discovery" category, because maturity varies substantially by modality and use case. A biotech evaluating a generative chemistry initiative for small-molecule leads is assessing a meaningfully more mature category than one evaluating de novo antibody generation, even though both fall under the same broad "generative AI in discovery" heading internally.

Regulatory Risk Tiering: The FDA and EMA Frameworks

Regulatory risk is not a peripheral consideration for biotech AI strategy; for any AI application touching nonclinical, clinical, postmarketing, or manufacturing phases of the drug product life cycle, it is often the dominant constraint. CDER (the FDA's Center for Drug Evaluation and Research) reports it has seen "a significant increase in the number of drug application submissions using AI components over the past few years," spanning nonclinical, clinical, postmarketing, and manufacturing phases (www.fda.gov). That growth trajectory is documented: CDER's January 2025 draft guidance was informed in part by the agency's experience with **over 500 submissions** with AI components between 2016 and 2023 (www.fda.gov).

That January 2025 draft guidance, titled *Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products*, establishes a **seven-step, risk-based framework** for model planning, development, validation, and monitoring (www.troutman.com). Legal analysis of the guidance notes it covers AI models used in "the nonclinical, clinical, post-marketing, and manufacturing phases of the drug product life cycle" where the model produces information to support a regulatory decision on safety, efficacy, or quality, and explicitly states the guidance "does not cover AI use in drug discovery or operational efficiencies that do not affect patient safety, drug quality, or study reliability" (www.goodwinlaw.com). This scoping detail matters for the prioritization matrix's regulatory risk axis: an early-discovery generative chemistry tool and a clinical-decision-support model can carry very different regulatory risk scores even though both are colloquially called "AI in drug development." The FDA sought public comment on the draft guidance through April 7, 2025, with particular interest in feedback on post-marketing pharmacovigilance use cases (www.goodwinlaw.com). The guidance provides recommendations to industry on the use of AI to produce information or data intended to support regulatory decision-making regarding the safety, effectiveness, or quality of drugs. CDER states it "plans to continue developing and adopting a risk-based regulatory framework that promotes innovation and protects patient safety" (www.fda.gov), and formed the CDER AI Council to consolidate governance across policy, regulatory, and technology efforts previously

spread across an AI Steering Committee, an AI Policy Working Group, and an AI Community of Practice (www.fda.gov). A regulatory attorney interviewed on the guidance framed its practical takeaways around "early FDA engagement and documentation expectations," underscoring that sponsors who wait until submission to think about AI credibility are already behind (www.troutman.com). Public comments on the guidance also raised concerns specific to generative and foundation models and to third-party AI, themes that will need to be resolved before the draft finalizes (www.troutman.com).

In January 2026, the FDA and EMA jointly published **ten Guiding Principles of Good AI Practice in Drug Development**, extending the seven-step framework's underlying philosophy into an internationally harmonized set of expectations spanning the full drug lifecycle, from early research through post-market safety (www.gxp-cc.com). Applied Clinical Trials Online reported that the principles "emphasize a human-centric, risk-based approach, focusing on data governance, multidisciplinary expertise, and transparent model development," and that they "are foundational, not prescriptive, and will evolve with technological advancements, prioritizing patient safety and ethical considerations" (www.appliedclinicaltrialsonline.com). The outlet also noted the initiative's broader aim: "The initiative supports global regulatory convergence and collaboration, aligning with ongoing efforts to integrate AI into regulatory frameworks" (www.appliedclinicaltrialsonline.com). Critically, the principles are not a checklist for approval but expectations regulators use to assess credibility, reliability, and fitness for purpose of AI/ML used across the drug development lifecycle, and they emphasize scaling controls proportionally to impact, allowing low-risk applications to progress with streamlined oversight while higher-risk uses undergo deeper validation and rigorous review (www.gxp-cc.com).

For biotech strategy teams, this proportionality principle is the operational bridge between the five-axis prioritization matrix and day-to-day governance. *Table 2* maps several of the ten principles to the functions responsible for implementing them.

Table 2. Selected FDA/EMA AI Principles Mapped to Organizational Roles

Principle	Key Roles Impacted	Practical Implication
Human-centric by design	Clinicians, Medical Affairs, Pharmacovigilance	Human experts retain final decision authority; AI is decision support, not a replacement for clinical judgment (www.gxp-cc.com)
Risk-based approach	Risk Management, Quality Assurance	Controls and oversight scale proportionally to potential impact, tied to patient safety and trial integrity, not the technology alone (www.gxp-cc.com)
Clear context of use	Regulatory Affairs, Program Leads	The precise intended use of each AI tool must be documented, including what it is not intended to do, anchoring risk assessment and validation (same source, principle 4)
Data governance and documentation	Data Management, Biostatistics	Data weaknesses are described as a leading cause of rejected AI evidence in regulatory submissions; full traceability and lineage are required (www.gxp-cc.com)
Adherence to standards	IT/Tech, Data Science	AI technologies must comply with relevant legal, ethical, technical, scientific, cybersecurity, and regulatory standards, including ICH guidelines and GxP (same source, principle 3)

Principle	Key Roles Impacted	Practical Implication
Multidisciplinary expertise	Cross-functional (All)	Clinical, statistical, computational, and regulatory expertise must be engaged across AI development and use, through formal governance structures with documented decision rights (same source, principle 5)
Model design and development	Data Scientists, ML Engineers	Reproducible, validated development pipelines must prioritize explainability, avoid overfitting, and justify model choices against the intended context of use (same source, principle 7)
Risk-based performance assessment	Validation/QA Teams	Testing must be scaled to risk, including across subgroups, sites, and stress scenarios, with performance criteria predefined rather than selected post hoc (same source, principle 8)
Life-cycle management	DevOps, Compliance	Ongoing governance including drift monitoring, version control, and planned decommissioning is required for the model's full lifecycle (same source, principle 9)
Clear, essential information	Regulatory, Communications	Documentation on AI technologies must be accessible and suited to the audience, regulators, reviewers, and end users, for submissions and assessments (same source, principle 10)

The regulatory environment described above should inform how organizations set the regulatory risk axis weight in the prioritization matrix discussed in the prior section. A commercial analytics use case, for example a next-best-action recommendation embedded in a customer relationship management (CRM) platform, sits at the low-risk end of this spectrum and can typically proceed with streamlined oversight. A model whose outputs support a regulatory decision on safety, effectiveness, or quality, for example evidence submitted in an investigational new drug (IND) application, sits at the high-risk end, where the FDA's draft guidance, which contains non-binding recommendations, recommends applying its seven-step credibility assessment with the greatest rigor. A generative model used only for early target selection falls outside that guidance's scope as noted above, although sponsors may still choose risk-tiered internal controls. Treating these as equivalent "AI projects" for governance purposes is one of the most common strategic errors organizations make, and it is precisely the error the risk-based, proportional-control principle is designed to prevent. Applied Clinical Trials Online notes that AI is already being used to generate and analyze evidence "across nonclinical, clinical, manufacturing, and post-marketing phases," and that while it has the potential to reduce development timelines and improve decision-making, its complexity requires careful oversight throughout its lifecycle (www.appliedclinicaltrialsonline.com).

Data Readiness as the Binding Constraint

If the prioritization matrix identifies which initiatives to fund, data readiness increasingly determines whether they can actually be delivered. Benchling's 2026 Biotech AI Report, drawing on a survey of approximately 100 biotechnology and pharmaceutical organizations conducted in November 2025, states plainly: "The reported number one reason AI pilots fail is due to challenges with data quality" (www.benchling.com). Only **6%** of surveyed organizations describe their R&D data environment as "fully integrated," with data curated, standardized, and accessible across R&D functions; **45%** describe their environment as "developing," with key datasets connected but inconsistent standards and gaps remaining.

ZS's CDIO survey corroborates this from the technology-leadership vantage point: **68%** of pharma and biotech leaders say neglecting data quality and governance early is the main reason AI initiatives fail, and CIOs are responding by increasing investment in cloud and infrastructure (**88%** of respondents), data products and platforms (**86%**), AI platforms (**84%**), and IT/DevOps tooling (**79%**) over the next 12 months. This is a strategy issue, not merely an information technology (IT) issue: the ZS survey found that AI experimentation is common but scaled deployment is rare, with only **40%** of pilots that clear an initial value threshold reaching scaled production. The report attributes the gap not purely to technology but to operating-model strain: technology and data capabilities (cited by **61%** of respondents), talent and skills (**58%**), and business engagement and decision-making (**56%**) are under the most pressure to change (www.zs.com).

Benchling's data further shows that AI adoption varies sharply by how validated and structured the underlying data is. High-adoption "killer app" use cases such as literature review (**76%** adoption), protein structure prediction (**71%**), scientific reporting (**66%**), and target identification (**58%**) "succeed due to clean, verifiable data that fits naturally into scientists' daily work" (www.benchling.com). By contrast, adoption drops in areas like generative design, biomarker analysis, and ADME (absorption, distribution, metabolism, and excretion), where data is scattered, incomplete, and hard to validate. ADME-specific adoption sits at only **29%**, versus **71%** for protein structure work. The practical takeaway for the prioritization matrix's data readiness axis is that it should be scored per use case, not per organization: a single biotech can simultaneously have high data readiness for literature synthesis and near-zero data readiness for biomarker-driven patient stratification.

Consistent with the data-readiness figures cited above, biotechs with high overall AI adoption report deeper wet-lab-to-dry-lab integration, with 30% describing that integration as "highly" connected, compared with just 18% among biotechs with low AI adoption, evidence that data integration and AI maturity reinforce each other rather than developing independently.

Talent sourcing compounds the data challenge. The same Benchling survey found the top source of AI talent is internal upskilling of existing scientific staff, at **67%**, well ahead of hiring from technology companies, at **21%**. This "build what differentiates, buy what scales" mindset, in the report's own phrase, has direct implications for the technical feasibility axis: organizations should not assume that hiring external AI talent alone resolves feasibility constraints when internal domain expertise is the actual bottleneck.

Sequencing: Fast-Track Wins Versus Long-Cycle Bets

Once individual use cases are scored, biotech and pharma leaders must sequence them across a portfolio, and ZS's research suggests budgets are effectively splitting along two distinct paths (www.zs.com). The fast track, where value is easier to prove on shorter cycles, includes enterprise tech and data operations, where **49%** of respondents already report consistently demonstrating measurable value, and commercial sales and marketing efforts such as healthcare provider engagement and personalization, where **47%** report the same. Supply chain and manufacturing lags at **29%** current impact but with **57%** expecting results within the next year.

The long game, discovery and clinical development, carries less certain but potentially larger payoffs. Only **17%** of respondents say they can prove measurable value today from data, digital, and AI investments in discovery, though **42%** expect value within the next year; clinical shows a similar pattern, with about **30%** demonstrating value now and **45%** anticipating progress within a year, particularly in study planning, patient recruiting, and real-world data use (www.zs.com). One pharma CIO quoted in the survey described this discipline directly: the organization has "learned enough with AI to move past experimentation" and is now concentrating on five to ten high-impact use cases, each with the potential to deliver 20 to 30% ROI.

This dual-path structure has a direct strategic implication: organizations should not force discovery and commercial AI initiatives through the same time-to-value bar. A discovery initiative scored on the same 12-month payback anchor used for a commercial CRM enhancement will systematically underscore against a criterion it was never designed to meet, understating its strategic value. ZS's own guidance to digital leaders is unambiguous on this point: nearly seven in ten respondents, **67%**, warn that launching an AI initiative without clear goals and success metrics defined up front is itself a mistake (www.zs.com).

This sequencing discipline extends to how digital and technology functions themselves are organized. ZS found that **61%** of pharma and biotech technology leaders plan to work with external partners to develop AI solutions, and **58%** describe their organization as taking a "trust-first" approach that embeds governance throughout the AI life cycle rather than bolting it on at the end (www.zs.com). CIOs are also actively reshaping their own operating models to support this discipline: **55%** report having the authority to reshape their enterprise operating model today, and **86%** are testing or making changes to roles and teams to ensure resources are deployed more effectively against the value agenda (www.zs.com).

At Sanofi, this dual-path thinking is embedded directly into portfolio governance. Helen Merianos, SVP and Global Head of R&D Strategy and Portfolio Management, described AI's role as helping the company decide "which programs to advance, accelerate, or stop," with AI recommendations weighing unmet need and likelihood of success against competitive differentiation risk (www.sanofi.com). Crucially, she also emphasized the human-in-the-loop discipline central to the FDA/EMA principles described above: "AI is not perfect, and it's only as good as the data it's trained on. That's why we have an organizational policy for responsible AI use, with humans always in the loop" (www.sanofi.com).

EY's analysis of AI-enabled portfolio decision-making reinforces why this governance layer matters at the portfolio office level specifically. The traditional R&D portfolio office role, EY notes, "has been to run simulations with dozens of complex inputs and then provide two- or three-dimensional visual aids and narrative detail to decision-makers," a process historically impaired by "poor-quality program data, long turnaround times for additional analysis by program teams, and analysis and visualization tools that require significant effort to update to answer new questions" (www.ey.com). AI's contribution, in EY's framing, is to elevate the portfolio office "to a strategic partnership with R&D and corporate leaders" through better inputs, digital twin and virtual pipeline analytics, and real-time decision support (www.ey.com). EY specifically highlights the "virtual pipeline" technique this enables: "a hypothetical model that tests various future assets against restraints, such as funding and personnel, to identify possible bottlenecks and facilitate strategic planning" (www.ey.com), which is functionally a scenario-testing layer on top of the five-axis prioritization matrix described earlier in this report.

Data Analysis and Evidence

This section consolidates the quantitative evidence underpinning biotech AI strategy decisions, spanning adoption rates, return-on-investment (ROI) benchmarks, market sizing, and capital flows.

Adoption by use case. Benchling's 2026 Biotech AI Report breaks adoption into scientific and non-scientific categories: **81%** of surveyed organizations use AI in highly specific scientific use cases such as biomarker discovery, molecule design, and experiment optimization, while **92%** use AI in non-scientific use cases such as document authoring, software engineering, and literature search (www.benchling.com). Within R&D specifically, adoption (defined as "using," not merely "piloting") breaks down as follows: target identification **58%**, protein structure and property prediction **71%**, biomarker identification and analysis **40%**, hit identification **37%**, lead optimization and selection **47%**, de novo generative drug design **42%**, synthetic biology and pathway design **42%**, experiment design and planning **39%**, ADME **29%**, and data analytics and visualization **50%** (www.benchling.com).

ROI and hit-rate benchmarks. AI-guided screening delivers hit rates of **22 to 46%**, compared with approximately **2%** for random screening, a 10 to 20-fold improvement, and multi-round AI optimization has been documented to yield 10-fold improvements in properties such as brain penetration (publisher's own synthesis, not an independent primary source). A peer-reviewed first analysis of AI-discovered molecules in clinical trials, published via PubMed, found an **80 to 90%** Phase I success rate for AI-discovered molecules, "substantially higher than historic industry averages," though the Phase II success rate of approximately **40%** was "comparable to historic industry averages," on a limited sample size (pubmed.ncbi.nlm.nih.gov). This is an important nuance often lost in vendor marketing: AI appears to materially de-risk the earliest, chemistry-driven stage of clinical development, but the harder, biology-driven Phase II efficacy hurdle has not yet shown the same improvement.

Market sizing. The global AI-in-drug-discovery market is valued at **$6.93 billion in 2025**, projected to reach **$7.62 billion in 2026** and approximately **$17.81 billion by 2035**, growing at a CAGR of **9.90%** from 2026 to 2035, according to Precedence Research (www.precedenceresearch.com). North America accounted for the largest market share, **56.18%**, in 2025, while the Asia-Pacific region is projected to grow fastest, at a **21.1%** CAGR from 2026 to 2035 (www.precedenceresearch.com). By therapeutic area, oncology accounted for **21%** of revenue share in 2025, and by application, the drug optimization and repurposing segment held around **51%** of 2024 revenue share, though the preclinical testing segment is growing faster (www.precedenceresearch.com). The U.S. market specifically is valued at **$3.15 billion** in 2026, projected to reach approximately **$7.50 billion** by 2035 at a 10.26% CAGR, reflecting the country's concentration of both AI-native biotech startups and large incumbent R&D budgets (www.precedenceresearch.com).

Capital flows and deal activity. As of early 2026, more than **173** AI-originated drug programs are reported in clinical development, up from roughly 24 in late 2023 (publisher's own synthesis, not an independent primary source). Notable deals include Xaira Therapeutics' **$1 billion** launch funding in April 2024, described by ARCH Venture Partners as "the largest initial funding commitment in ARCH history" (www.fiercebiotech.com); Recursion's **$688 million** all-stock acquisition of Exscientia, announced in August 2024, under which Exscientia shareholders received 0.7729 Recursion shares per Exscientia share and which was expected to leave the combined entity with about **$850 million** in cash to fund roughly three years of operations

(www.reuters.com); and Eli Lilly's collaboration with Insilico Medicine, under which Insilico is eligible for a **$115 million** upfront payment plus development, regulatory, and commercial milestones that could bring total deal value to approximately **$2.75 billion**, plus tiered royalties (insilico.com).

Scaling and abandonment risk. Only **22%** of life sciences leaders report having successfully scaled AI, and just **9%** report achieving significant returns, per Deloitte's 2026 survey (www.deloitte.com). Gartner's earlier prediction, discussed above, was that at least 30% of generative AI projects would be abandoned after proof of concept by the end of 2025 due to poor data quality, inadequate risk controls, escalating costs, or unclear business value. *Table 3* summarizes these adoption and value-realization statistics side by side.

Table 3. Key AI Adoption and Value-Realization Statistics in Biotech and Pharma (as of July 2026)

Metric	Value	Source
Life sciences leaders who have successfully scaled AI	22%	Deloitte 2026 Life Sciences Outlook (cited above)
Life sciences leaders reporting significant AI returns	9%	Deloitte 2026 Life Sciences Outlook (cited above)
AI pilots that reach scaled deployment (of those with clear value)	40%	ZS 2026 CDIO Outlook Survey (cited above)
Leaders citing data quality/governance neglect as main failure cause	68%	ZS 2026 CDIO Outlook Survey (cited above)
AI-guided screening hit rate (versus ~2% random)	22 to 46%	IntuitionLabs synthesis of published hit-rate studies (intuitionlabs.ai)
Phase I success rate for AI-discovered molecules	80 to 90%	Peer-reviewed study, PubMed (pubmed.ncbi.nlm.nih.gov)
Global AI-in-drug-discovery market size, 2026	$7.62 billion	Precedence Research (cited above)
Global AI-in-drug-discovery market CAGR, 2026 to 2035	9.90%	Precedence Research (cited above)

The pattern across every row of the table above is consistent: adoption and enthusiasm figures (percentages using AI in at least one workflow) are consistently far higher than value-realization figures (percentages reporting successful scaling or significant returns). This adoption-to-realization gap is the single most important quantitative finding for biotech AI strategy: it is the empirical justification for treating prioritization, data readiness, and governance as first-order strategic concerns rather than downstream implementation details.

Case Studies and Real-World Examples

Insilico Medicine: Generative Chemistry to Global Pharma Partnership

Insilico Medicine, a Hong Kong Exchange-listed clinical-stage biotechnology company (HKEX: 3696), has used generative AI and automation since 2014 to identify novel therapeutic targets, design new molecules, and select candidates for testing through its Pharma.AI platform (aws.amazon.com). Using Amazon SageMaker,

the company achieved an **83%** reduction in time to deploy model updates and more than **16 times** accelerated model iteration and deployment time (aws.amazon.com). The company is noted as one of the first AI-driven drug discovery companies to nominate a preclinical candidate in under 18 months and reach human clinical trials in 30 months (aws.amazon.com). In 2026, Insilico announced a collaboration with Eli Lilly that grants Lilly an exclusive worldwide license for a portfolio of preclinical oral therapeutics, with Insilico eligible for a **$115 million** upfront payment and total deal value of approximately **$2.75 billion** including milestones, plus tiered royalties (insilico.com). Founder and CEO Alex Zhavoronkov framed the collaboration around Insilico's platform breadth: "By deploying frontier AI technologies that scale from biomarkers to life models, world models of human and animal life, we can identify multi-purpose targets driving multiple diseases at the same time" (insilico.com). The deal structure itself is instructive for the time-to-value axis of the prioritization matrix: Insilico disclosed the arrangement carries a **$115 million** upfront payment with the remainder contingent on development, regulatory, and commercial milestones, meaning the headline **$2.75 billion** figure represents a ceiling contingent on years of further clinical execution, not cash received at signing (insilico.com).

Recursion Pharmaceuticals and Exscientia: Consolidation as Strategy

In August 2024, Recursion Pharmaceuticals, a Salt Lake City-based biotech using AI to discover new drug candidates, agreed to acquire smaller rival Exscientia in an all-stock deal valued at **$688 million** (www.reuters.com). Under the deal terms, Exscientia shareholders received 0.7729 Recursion shares for each Exscientia share held, and the combined entity was expected to have about **$850 million** in cash to fund operations for approximately three years (www.reuters.com). Reuters reported the deal came "as major drugmakers double down on AI to boost drug development, find patients for clinical trials quickly or reduce the number of people needed to test medicines, both accelerating drug development and potentially saving millions of dollars" (www.reuters.com). The combined company retained the Recursion name under CEO Chris Gibson, with Exscientia's interim CEO David Hallett becoming science chief, and STAT News reported that Recursion shareholders retained approximately **74%** ownership of the combined company while Exscientia shareholders held approximately **26%** (www.statnews.com). Exscientia's approach prior to the merger combined generative AI algorithms for in silico compound design with automated robotic synthesis in what it called the design-make-test-learn (DMTL) cycle, built on Amazon Web Services, with the company framing the goal as solving "for efficiency and effectiveness" against a conventional discovery process that "can take up to 15 years and cost over 2 billion dollars, with an average failure rate of 90 to 96 percent" (aws.amazon.com). Prior to the merger, Exscientia stated it had "accelerated drug design by up to 70 percent while decreasing capital cost by 80 percent, compared with industry benchmarks," using a synthesis-aware, iterative approach that made 10 times fewer compounds than the industry average (aws.amazon.com). Post-merger, Recursion has continued to publicize AI-designed programs such as REC-1245, which Recursion reports as progressing from novel target biology to preclinical drug candidate in under 18 months, described by Recursion CEO Chris Gibson as "nearly twice the speed of the industry average" (www.genengnews.com).

Xaira Therapeutics: A Billion-Dollar AI-Native Launch

Xaira Therapeutics launched in April 2024 with **$1 billion** in committed funding, incubated by ARCH Venture Partners and Foresite Labs, with Marc Tessier-Lavigne, former chief scientific officer at Genentech, at the helm (www.fiercebiotech.com). The financing included F-Prime, NEA, Sequoia Capital, Lux Capital, and Lightspeed Venture Partners among others, and ARCH managing director Robert Nelsen called it "the largest initial funding commitment in ARCH history" (www.fiercebiotech.com). Xaira was co-founded by David Baker, director of the Institute for Protein Design at the University of Washington, and combines machine learning, data generation, and therapeutic product development, incorporating researchers behind the RFdiffusion and RFantibody protein and antibody design models (www.fiercebiotech.com). Tessier-Lavigne framed the strategic bet in terms directly relevant to prioritization discipline: "Now, witnessing how AI is impacting other industries and the considerable progress in applications of AI in biology, I believe we are poised for a revolution" (www.fiercebiotech.com). Xaira illustrates a distinct strategic archetype from Sanofi and Pfizer: rather than integrating AI into an existing large-pharma R&D organization, it was built AI-native from inception, with prioritization decisions embedded in its founding scientific thesis around protein and antibody design.

Sanofi: AI-Governed Portfolio Decision-Making at Enterprise Scale

Sanofi represents the large-pharma pattern of embedding AI directly into R&D portfolio governance rather than treating it as a discrete discovery tool. Helen Merianos, SVP and Global Head of R&D Strategy and Portfolio Management, described AI's role in making decisions about which programs to advance, accelerate, or stop, and described the shift as company-wide transformation: "I'm amazed at the progress we've made in portfolio management by leveraging AI as a key accelerator" (www.sanofi.com). Sanofi's AI applications span molecule design through the writing of clinical study reports, including using AI and in silico modeling before investigational compounds enter human trials, which the company states "raises the likelihood of success in the clinic" (www.sanofi.com). Notably, Merianos emphasized that governance and human oversight are not optional add-ons but load-bearing elements of the strategy: "Even if the AI isn't 100% right, it still makes us smarter. But without human perspectives, governance committees may lose trust. Real people remain essential to interpret nuance, recognize blind spots, and ensure confidence" (www.sanofi.com).

Pfizer: Embedding AI Across an Existing Commercial Pipeline

Pfizer's development of the COVID-19 oral antiviral Paxlovid illustrates how a large incumbent pharmaceutical company can embed AI into an already-running commercial pipeline rather than building a standalone AI unit. Pfizer has stated that supercomputing capabilities paired with advanced computational methods helped its scientists optimize their search for molecules that could deliver Paxlovid orally rather than intravenously; this is a company statement that has not been independently verified (pfizer.com). Pfizer is a useful counterpoint to Xaira in this report's case studies: rather than a single-thesis platform bet, it represents AI adoption threaded incrementally through an existing, highly regulated commercial and clinical operation, where the five-axis prioritization matrix's regulatory risk and technical feasibility axes are shaped by pre-existing systems and submission history rather than a greenfield build.

Implications and Future Directions

Several structural trends will shape biotech AI strategy over the next 12 to 24 months. First, the regulatory environment will continue converging internationally. The joint FDA/EMA ten-principles document, published in January 2026 and discussed above, follows the 2024 FDA-EU bilateral meeting and builds on the EMA's own July 2023 draft reflection paper, which outlined "principles relevant to the application of AI and machine learning (ML) at any step of a medicines' lifecycle, from drug discovery to the post-authorisation setting" and likewise emphasized that "a human-centric approach should guide all development and deployment of AI and ML" (www.ema.europa.eu), and the FDA's original January 2025 draft guidance is listed on the FDA's guidance page as a draft ("Not for implementation"), and was informed by more than 800 external comments received on a May 2023 discussion paper and a public workshop held in August 2024, with a further public workshop held in October 2025 (www.fda.gov). Organizations that have already mapped their AI portfolio against the seven-step credibility framework and the ten guiding principles will be positioned to move faster once the guidance finalizes; those treating regulatory alignment as a downstream compliance task rather than an upfront prioritization criterion will face costly retrofitting. The EMA's reflection paper is itself part of a broader multi-year workplan between European regulators on data and AI governance, signaling that the convergence trend documented here is structural rather than a one-off coordination exercise between two agencies (www.ema.europa.eu).

Second, the adoption-to-realization gap documented throughout this report, 22% scaled versus far higher pilot rates, 40% of qualifying pilots reaching production, is unlikely to close through additional AI tooling alone. Gabriele Ricci, Takeda's Chief Data and Technology Officer, quoted in Deloitte's 2026 outlook, captured the industry mood directly: "We're all entering a period of purposeful transformation, where discipline and innovation must coexist as the industry matures beyond hype toward measurable productivity from AI and data" (www.deloitte.com). This suggests the next competitive differentiator in biotech AI strategy will be organizational discipline, prioritization rigor, data governance, and lifecycle monitoring, rather than access to more capable models, which are becoming commoditized.

Third, agentic AI, cited by **30%** of Deloitte's 2026 respondents as an influential trend, a new survey category this year (www.deloitte.com), and by 41 to 45% of ZS's respondents as a near-term focus area for IT operations and R&D discovery workflows respectively (detailed in the FAQ below), will require the prioritization matrix's regulatory risk axis to be applied with particular care. Autonomous agents that can act, decide, and execute tasks introduce a categorically different risk profile than task-level copilots, and ZS notes that in customer- and patient-facing functions, most executives remain cautious, favoring smaller-scale use cases with human-in-the-loop checks over full autonomy.

Fourth, commercial technology platforms will play a growing role in translating strategic AI priorities into deployed capability. Veeva's Vault AI, embedded across its Vault Platform and applications, illustrates how AI capability is increasingly delivered through existing systems of record rather than as standalone tools, with the company positioning it as "deep AI that's simple, secure, and compliant for life sciences companies of all sizes" (www.veeva.com). For organizations whose commercial and quality systems already run on Vault CRM, prioritization decisions around commercial AI use cases increasingly become platform configuration and extension decisions (dashboards, X-Pages applications, CRM extensions) rather than build-versus-buy

decisions from scratch, which shifts technical feasibility scoring meaningfully in favor of use cases that extend the existing platform footprint. Organizations evaluating this path benefit from working with implementation partners who understand both the underlying compliance requirements, including 21 CFR Part 11 and EU Annex 11 audit trail obligations, and the specific capabilities of Vault CRM X-Pages, the same product as MyInsights built for the Vault Platform, which is included in the standard Vault CRM license at no additional cost (www.veeva.com).

Finally, the market data reviewed above (a projected 9.90% CAGR through 2035, more than 173 AI-originated programs in clinical development, and successive multi-hundred-million to multi-billion-dollar deals) indicates capital markets remain willing to fund AI-native biotech theses even amid uneven scaling results at incumbent pharma companies. This divergence, well-funded pure-play AI biotechs on one hand and struggling-to-scale incumbent AI programs on the other, is itself a strategic signal: prioritization frameworks calibrated for a resource-constrained internal portfolio (the incumbent case) look different from those calibrated for a single-thesis, fully-funded platform company (the Xaira or Insilico case), and organizations should be explicit about which archetype they are in before importing a framework wholesale from a peer of the other type.

Frequently Asked Questions (FAQs)

What is biotech AI strategy? Biotech AI strategy is the discipline of selecting, sequencing, funding, and governing artificial intelligence initiatives across a biotechnology or pharmaceutical organization's research, clinical, manufacturing, and commercial functions, using structured criteria such as strategic value, technical feasibility, data readiness, regulatory risk, and time-to-value rather than ad hoc or purely technology-driven decision-making, as described in the prioritization framework discussed earlier in this report.

How do you prioritize AI projects in a biotech organization? The most defensible approach scores each candidate use case against multiple weighted axes, typically strategic value, technical feasibility, data readiness, regulatory risk, and time-to-value, using anchored 1-to-5 scoring definitions and a calibration session with a cross-functional governance group before scoring begins, then converts the weighted total into a tier (A through D) to support portfolio decisions, per the five-axis matrix detailed above.

What are the most common biotech AI use cases today? As of the most recent industry survey data cited earlier in this report, the highest-adoption use cases are literature and knowledge extraction (76% adoption), protein structure and property prediction (71%), scientific reporting (66%), and target identification (58%), because these areas benefit from cleaner, more verifiable data than areas such as generative molecular design or biomarker analysis.

What is the ROI of AI in drug discovery? ROI should be measured across financial, operational, clinical, and scientific-outcome dimensions rather than time savings alone. Concrete benchmarks include AI-guided screening hit rates of 22 to 46% versus roughly 2% for random screening, and an 80 to 90% Phase I clinical trial success rate for AI-discovered molecules, though the Phase II rate of about 40% remains comparable to historic industry averages (the 22 to 46% hit-rate figure is the publisher's own synthesis, not an independent primary source) (pubmed.ncbi.nlm.nih.gov).

What regulatory framework governs AI use in drug development? In the United States, the FDA's January 2025 draft guidance establishes a seven-step, risk-based credibility assessment framework for AI models supporting regulatory submissions (www.troutman.com). Internationally, the FDA and EMA jointly published ten Guiding Principles of Good AI Practice in Drug Development in January 2026, which scale oversight proportionally to risk across the full drug lifecycle, as detailed earlier in this report.

Why do most biotech AI pilots fail to scale? The primary cited reason is data quality and governance, not model capability. ZS's 2026 CDIO survey, cited earlier in this report, found 68% of pharma and biotech leaders cite neglecting data quality and governance early as the main reason AI initiatives fail, and only 40% of pilots that clear an initial value bar reach scaled deployment.

How big is the AI in drug discovery market? Precedence Research, cited earlier in this report, values the global AI-in-drug-discovery market at approximately $7.62 billion in 2026, projected to reach $17.81 billion by 2035 at a 9.90% CAGR.

Should a small or mid-cap biotech use the same AI prioritization approach as a large pharma company? The underlying five-axis logic (strategic value, feasibility, data readiness, regulatory risk, time-to-value) applies broadly, but weighting should differ. Smaller organizations with acute data infrastructure gaps often need to weight data readiness more heavily, around 25% or higher, since data preparation work can otherwise consume a use case's entire budget before value is realized, consistent with the weighting adjustments discussed earlier in this report.

How does agentic AI change biotech AI prioritization? Agentic AI, systems that can plan, act, and adapt with limited human intervention, is moving fastest in IT operations, where **45%** of surveyed pharma and biotech leaders plan to build agentic workflows, and in R&D discovery, where **41%** have similar plans, but customer- and patient-facing functions remain more cautious, favoring human-in-the-loop task automation over full autonomy (www.zs.com). Because autonomous, multi-step systems are harder to validate against FDA and EMA risk-based credibility expectations, agentic initiatives should generally be scored with a lower technical feasibility score and a higher regulatory risk score than single-step copilot tools until validation practices mature.

Should biotech AI talent be hired externally or developed internally? Benchling's survey, cited earlier in this report, found the leading source of AI talent is internal upskilling of existing scientific staff, at 67%, ahead of hiring from AI-focused biotechnology companies (39%), academic institutions (31%), and technology companies (21%). This favors a "build what differentiates, buy what scales" approach over wholesale external hiring or wholesale outsourcing.

Conclusion

Biotech AI strategy in 2026 is defined less by a shortage of viable use cases than by an excess of them relative to organizational capacity to execute, govern, and scale. The evidence assembled in this report points to a consistent set of conclusions. Structured, multi-axis prioritization, weighing strategic value, technical feasibility, data readiness, regulatory risk, and time-to-value, converts an unmanageable list of candidate initiatives into a defensible, board-ready sequence, and calibration discipline is what keeps that structure honest rather than performative. Data readiness, more than algorithmic sophistication, determines whether a

well-scored initiative actually reaches production, a pattern visible across every survey examined in this report, from Benchling's use-case-level adoption breakdown to ZS's operating-model pressure findings. Regulatory risk is not a separate workstream bolted onto AI strategy after the fact; the FDA's draft, non-binding seven-step credibility framework and the FDA/EMA's ten guiding principles both point toward proportional, risk-tiered oversight as the operating model AI initiatives are best designed around from the outset, not retrofitted into later.

The gap between AI adoption and AI value realization, 22% of life sciences leaders reporting successful scaling against far higher piloting and experimentation rates, is the central strategic fact organizations must confront. Closing that gap requires treating prioritization as a recurring governance discipline rather than a one-time exercise: use cases should be rescored as data infrastructure matures, as regulatory guidance finalizes, and as organizational capacity to execute changes. The named cases examined here, spanning AI-native platform companies funded on single scientific theses, consolidating mid-cap AI biotechs, and large incumbent pharma companies embedding AI into existing portfolio governance, demonstrate that there is no single correct organizational archetype for biotech AI strategy, but there is a common discipline underlying the ones that are working: explicit criteria, calibrated scoring, data-readiness honesty, and risk-proportional governance, sustained over multiple review cycles rather than applied once at initiative launch.

For life sciences organizations evaluating how to structure this discipline internally, particularly where AI initiatives touch existing commercial and quality systems such as Veeva Vault CRM, the practical starting point is an honest inventory: which candidate initiatives exist, how they score against the five axes, what data infrastructure gaps stand between the current state and each initiative's data readiness requirement, and which regulatory risk tier each initiative falls into under the FDA/EMA proportionality principle. That inventory, revisited quarterly rather than built once and shelved, is what separates the roughly one in five organizations reporting successful AI scaling from the majority still experimenting.

References

- Deloitte, *2026 Life Sciences Outlook* (www.deloitte.com)
- Gartner, *Lack of AI-Ready Data Puts AI Projects at Risk* (www.gartner.com)
- U.S. Food and Drug Administration, *Artificial Intelligence for Drug Development* (CDER) (www.fda.gov)
- U.S. Food and Drug Administration, *Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products*, draft guidance, January 2025 (www.fda.gov)
- European Medicines Agency, *Reflection Paper on the Use of Artificial Intelligence in the Lifecycle of Medicines* (www.ema.europa.eu)
- Applied Clinical Trials Online, *FDA and EMA Align on Ten Principles to Guide Artificial Intelligence Use in Drug Development* (www.appliedclinicaltrialsonline.com)
- Troutman Pepper Locke, *Navigating FDA's 2025 AI Guidance: Risk-Based Framework, Public Comments, and Generative Models* (www.troutman.com)
- Goodwin Law, *FDA Publishes Its First Draft Guidance on Use of Artificial Intelligence in the Development of Drugs and Biological Products* (www.goodwinlaw.com)
- ZS, *Scaling AI in Pharma and Biotech: 2026 CDIO Research* (www.zs.com)

- pharmaphorum, *A Strategic Framework for Evaluating AI in Drug Discovery* (pharmaphorum.com)
- Sanofi, *AI Across the R&D Value Chain: Portfolio Decision-Making* (www.sanofi.com)
- EY, *How AI Can Support Drug Development Portfolio Decisions* (www.ey.com)
- Precedence Research, *Artificial Intelligence in Drug Discovery Market* (www.precedenceresearch.com)
- PubMed, *How Successful Are AI-Discovered Drugs in Clinical Trials? A First Analysis and Emerging Lessons* (pubmed.ncbi.nlm.nih.gov)
- IntuitionLabs, *Pharma AI Strategy: Scaling Digital Transformation*
- IntuitionLabs, *Veeva X-Pages Partner: Vault CRM and Full Ecosystem*
- FierceBiotech, *New AI Drug Discovery Powerhouse Xaira Rises With $1B in Funding* (www.fiercebiotech.com)
- Insilico Medicine, *Insilico Medicine Announces Global R&D Collaboration with Lilly* (insilico.com)
- Reuters, *Biotech Firm Recursion to Buy Smaller Peer Exscientia for $688 Million* (www.reuters.com)
- STAT News, *AI Drug Developers Recursion, Exscientia to Merge* (www.statnews.com)
- Amazon Web Services, *Insilico Medicine Accelerates Drug Discovery Using Amazon SageMaker* (aws.amazon.com)
- Amazon Web Services, *Exscientia Uses Generative AI to Reimagine Drug Discovery* (aws.amazon.com)
- Genetic Engineering and Biotechnology News, *Recursion's Fast-Track Road to Therapeutics Using AI-Based Maps of Biology* (www.genengnews.com)
- Veeva Systems, *Vault AI* (www.veeva.com)

Source links

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

1. https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html
2. https://www.zs.com/insights/scaling-ai-in-pharma-cdio-2026
3. https://pharmaphorum.com/deep-dive/strategic-framework-evaluating-ai-drug-discovery
4. https://www.benchling.com/biotech-ai-report-2026
5. https://pubmed.ncbi.nlm.nih.gov/38692505
6. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological
7. https://www.reuters.com/markets/deals/biotech-firm-recursion-buy-smaller-peer-exscientia-688-million-2024-08-08
8. https://insilico.com/news/uiy12zcjg1-insilico-medicine-announces-global-rampd
9. https://www.precedenceresearch.com/artificial-intelligence-in-drug-discovery-market
10. https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk
11. https://sakaradigital.com/blog/ai-use-case-prioritization-matrix-clients
12. https://www.gartner.com/en/newsroom/press-releases/2024-07-29-gartner-predicts-30-percent-of-generative-ai-projects-will-be-abandoned-after-proof-of-concept-by-end-of-2025
13. https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development
14. https://www.troutman.com/insights/navigating-fdas-2025-ai-guidance-risk-based-framework-public-comments-and-generative-models

15. https://www.goodwinlaw.com/en/insights/publications/2025/01/alerts-lifesciences-aiml-fda-publishes-its-first-draft-guidance
16. https://www.gxp-cc.com/insights/blog/fda-ema-ai-principles-role-based-roadmap-for-regulated-pharma-companies
17. https://www.appliedclinicaltrialsonline.com/view/fda-ema-align-ten-principles-artificial-intelligence-use-drug-development
18. https://www.sanofi.com/en/magazine/ai-in-healthcare/ai-across-the-research-development-value-chain-portfolio-decision-making
19. https://www.ey.com/en_us/insights/life-sciences/how-ai-can-support-drug-development-portfolio-decisions
20. https://www.fiercebiotech.com/biotech/new-ai-drug-discovery-powerhouse-xaira-rises-1b-funding
21. https://aws.amazon.com/solutions/case-studies/insilico-customer-case-study
22. https://www.statnews.com/2024/08/08/recursion-exscientia-ai-drug-development-merger
23. https://aws.amazon.com/solutions/case-studies/exscientia-generative-ai
24. https://www.genengnews.com/topics/artificial-intelligence/recursions-fast-track-road-to-therapeutics-using-ai-based-maps-of-biology
25. https://www.pfizer.com/sites/default/files/investors/financial_reports/annual_reports/2022/story/data-and-ai-are-helping-to-get-medicines-to-patients-faster
26. https://www.ema.europa.eu/en/news/reflection-paper-use-artificial-intelligence-lifecycle-medicines
27. https://www.veeva.com/products/vault-ai
28. https://www.veeva.com/products/crm-suite/x-pages

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.