How to Build an AI Roadmap for Biotech Companies in 2026

7/1/2026 • 45 min read

ai roadmap for biotech | biotech ai strategy | ai in life sciences consulting | artificial intelligence in drug discovery | ai governance in life sciences | biopharma ai use cases | pharma ai adoption | life sciences digital transformation

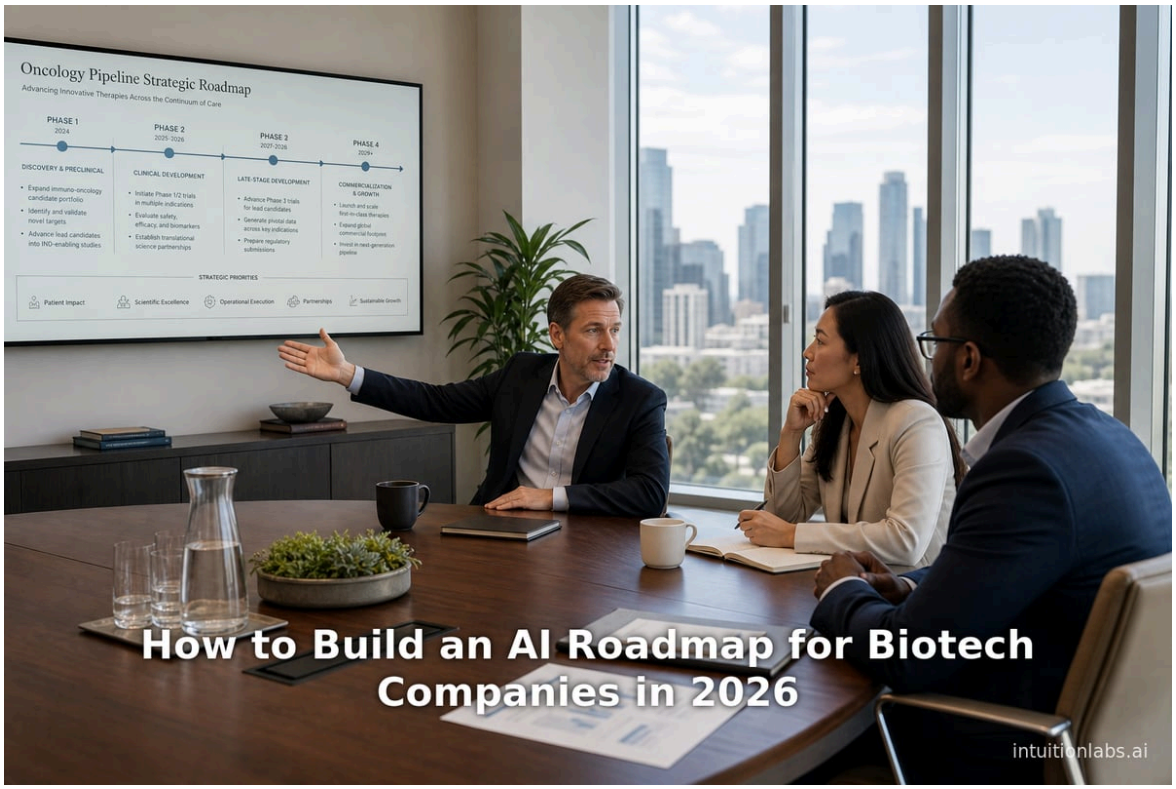

Executive Summary

Building an artificial intelligence (AI) roadmap for a biotech company in 2026 means sequencing a small number of high-value use cases against a realistic data, talent, and governance foundation rather than chasing every available pilot. The urgency is real: Global Market Insights valued the global AI in drug discovery market at $3.1 billion in 2025 and projects growth to **$43.9 billion by 2035**, a 30.5% compound annual growth rate (CAGR) ([1] gminsights.com), while Roots Analysis pegs the 2026 market at **$8.6 billion**, rising to $25.0 billion by 2035 at a 12.6% CAGR ([2] rootsanalysis.com). The wide spread between forecasters, also visible in IMARC Group's $2.2 billion 2025 baseline ([3] imarcgroup.com), signals a market still being defined, which is itself a planning input: biotech leaders should treat any single market number as directional, not a budget anchor. Deloitte frames the opportunity differently still, estimating that generative AI specifically could **"unlock $5 billion to $7 billion in value"** across research, manufacturing, and commercial functions, a figure it derives from research across 20 specific AI use cases rather than a broad market-sizing exercise ([4] deloitte.com).

The evidence on execution is sobering. McKinsey's late-2024 survey of more than 100 pharma and medtech leaders found that while all respondents had experimented with generative AI, only **32% had taken steps to scale it, and just 5% reported realizing it as a consistent, significant financial differentiator** ([5] mckinsey.com). That pattern echoes MIT's 2025 "State of AI in Business" research, which found that roughly **95% of generative AI pilots across industries fail to move the needle on profit and loss**, largely because of a "learning gap" in enterprise integration rather than model quality ([6] fortune.com). Boston Consulting Group's biopharma-specific research is more encouraging for the minority who commit: companies that adopt an "AI-first" operating model, following what BCG calls the **10/20/70 principle** (10% algorithms, 20% data and technology, 70% people and process), have compressed early discovery and candidate identification from four or five years to **eight months** and report a **5% to 15% revenue uplift** ([7] bcg.com ([8] bcg.com).

A defensible roadmap therefore rests on five sequenced pillars: (1) executive and board-level strategic alignment, given that only **8% of Russell 3000 and S&P 500 companies currently disclose board-level AI oversight** ([9] iss-stoxx.com); (2) data infrastructure readiness, since fragmented, non-interoperable lab and clinical data remains the top-cited obstacle; (3) disciplined use-case selection and sequencing across discovery, clinical development, manufacturing, and commercial operations; (4) governance aligned to the FDA's risk-based credibility framework and the January 2026 FDA-EMA joint principles (ema.europa.eu) as well as the European Union's Artificial Intelligence Act (Regulation (EU) 2024/1689), whose high-risk obligations become enforceable from 2 August 2026 (digital-strategy.ec.europa.eu); and (5) talent and change management, addressing a skills gap that the Pistoia Alliance found **34% of life-science R&D professionals now cite as a barrier to AI adoption, up from 23% in 2024** ([10] pistoiaalliance.org).

Real-world deployments illustrate both the promise and the risk of moving too fast or too slow. Novo Nordisk's 2026 commitment to embed OpenAI's tools across its entire pipeline, Roche's 3,500-plus Nvidia Blackwell graphics processing unit (GPU) "AI factory," and Insilico Medicine's advance of its AI-discovered drug rentosertib into Phase III trials for idiopathic pulmonary fibrosis all show what disciplined scaling can achieve ([11] roche.com ([12] clinicalresearchnewsonline.com). Conversely, BenevolentAI's repeated rounds of layoffs and pivot back to a "TechBio" model, and Recursion's post-merger pipeline cuts after acquiring Exscientia, show that a roadmap without financial discipline or a validated operating model can unravel quickly ([13] fiercebiotech.com). As an adjacent life-sciences and AI consultancy, and a Veeva Systems ecosystem partner, IntuitionLabs' own analysis of the pharmaceutical AI skills gap concludes that reskilling existing staff can be roughly half the cost of external hiring while improving retention ([14] intuitionlabs.ai), a finding this report treats as a central input for the talent section of any biotech AI roadmap.

Introduction and Background

A biotech AI roadmap is the sequenced plan by which a biopharmaceutical, medical-device, or life-sciences technology company moves from isolated experimentation with artificial intelligence to durable, governed, value-generating deployment across its research and development (R&D), clinical, manufacturing, and commercial functions. As of July 2026, the pressure to have such a plan has intensified sharply. Large pharmaceutical companies are no longer running one-off proofs of concept: Novo Nordisk announced it is embedding OpenAI's tools across "every stage of drug development, from discovery to manufacturing and sales," with **full integration targeted by the end of 2026** ([15] nexchron.com).Roche has built what it calls the industry's largest announced hybrid-cloud AI infrastructure, embedding Nvidia's BioNeMo platform into its "Lab-in-the-Loop" discovery process ([16] roche.com). For smaller and mid-sized biotechs without Roche's balance sheet, the strategic question is not whether to adopt AI but how to sequence a limited set of investments for maximum probability of return.

This shift has been building since well before generative AI captured headlines. Deloitte's early life-sciences AI survey found that **more than 60% of life sciences companies spent over $20 million on AI initiatives in 2019**, with process efficiency (43%) the most commonly realized benefit ([17] deloitte.com). What changed after 2023 was the arrival of large language models and foundation models capable of generalizing across biological modalities, which McKinsey's July 2023 analysis estimated could unlock **$60 billion to $110 billion a year in economic value across the pharmaceutical and medical-products industry's value chain** ([18] mckinsey.com). That estimate has propelled boardroom interest, but the gap between ambition and realized value, documented by McKinsey's own 2024 survey and MIT's 2025 research on enterprise generative AI, is exactly what a roadmap must close.

The scientific case for AI's contribution to drug discovery has also strengthened alongside the commercial one. Isomorphic Labs, the digital-biology company spun out of Google DeepMind, was built on the success of AlphaFold, work that Nature named its 2021 "Method of the Year" and Science its "Breakthrough of the Year," before Isomorphic extended that structural-prediction capability into a broader AI drug design engine covering small molecules and nucleic acids ([19] prnewswire.com). That lineage matters for a roadmap because it illustrates a recurring pattern in life-sciences AI: today's production tools, whether Isomorphic's design engine, Insilico Medicine's Pharma.AI platform, or Owkin's K Pro agents, typically trace back to a narrower academic or research breakthrough that took years to mature into a validated, regulator-facing capability. A roadmap that expects an equally fast maturation curve for a brand-new AI technique is likely to be disappointed; a roadmap that budgets multi-year validation cycles for genuinely novel AI methods, while moving faster on already-validated ones, is better calibrated to the industry's actual track record.

A useful roadmap is distinct from a technology inventory or a vendor shortlist. It sets sequencing logic (which use cases go first and why), assigns accountability (who owns data quality, who owns model risk, who owns the board relationship), and defines the gates a project must clear before scaling from pilot to production. Within the broader drug discovery market, of which AI is one growing slice, Mordor Intelligence separately reports that **"high-throughput screening delivered 34.27% of technology revenue in 2025,"** even as AI-specific platforms grow fastest of any technology segment tracked, evidence that AI is displacing incumbent screening methods gradually rather than all at once ([20] mordorintelligence.com). This report builds that roadmap in five parts: strategic foundations and governance; data infrastructure and readiness; use-case selection and sequencing across the value chain; regulatory and risk management; and talent and organizational change. It then examines the quantitative evidence base, profiles five named real-world cases, one of them a cautionary tale, and closes with implications for the next 12 to 24 months. Throughout, the report distinguishes vendor claims from independently measured outcomes and flags where credible sources disagree, since a roadmap built on an inflated number is a roadmap that will disappoint its sponsors.

Key Changes in How Biotechs Are Building AI Roadmaps

Establishing Strategic Foundations and Executive Alignment

The first and most frequently skipped step in an AI roadmap is deciding, at the executive and board level, what problem AI is actually meant to solve. BCG's 2026 analysis of AI-first biopharma companies found that the differentiator between organizations extracting value and those stuck in "pilot purgatory" was not algorithmic sophistication but organizational commitment: AI-first companies "focus on fewer, higher-value use cases and scale them across the enterprise, rather than scattering resources across pilots," backed by long-term capital commitments and a higher share of revenue dedicated to digital and AI initiatives than peers ([21] bcg.com). BCG's central heuristic, the 10/20/70 principle, holds that only 10% of the AI adoption challenge is algorithmic; 20% is data and technology; and a full **70% is people, process, culture, governance, and change management** ([22] bcg.com). This is the single most important planning fact for a biotech board deck on AI strategy: the budget line item for compute and licenses should not be mistaken for the budget line item that actually determines success.

Board-level accountability remains the weakest link industry-wide. Governance data firm ISS-STOXX reviewed 3,048 Russell 3000 and S&P 500 companies and found that only **245 companies (8%) disclosed board-level oversight of AI**, and only **275 (9%) had established formal AI policies**, while just **481 companies (16%) reported a director with specialized AI skills** ([9] iss-stoxx.com). Health Care is one of the five sectors where oversight and skills concentrate, but even there ISS-STOXX describes a "lone expert" pattern, where a single director carries AI fluency for the entire board, creating **"key-person risk"** ([23] iss-stoxx.com). A defensible roadmap therefore starts with a written charter naming an executive sponsor (frequently a chief digital or chief technology officer, a role several large pharmaceutical companies have already created specifically to own enterprise AI infrastructure), a board committee or designated director with AI literacy, and a small number (typically three to five) of enterprise-level objectives against which every AI investment is scored.

Deloitte's framing of this inflection point remains apt years after it was first published: AI is moving "from a 'nice to have' to a 'must have,'" which means companies "should build a vision and strategy to leverage AI, then put in place the building blocks needed to scale its use" ([24] deloitte.com). The ordering in that sentence, vision and strategy before scaling infrastructure, is precisely the sequence this report's five pillars follow, and it is also the ordering most frequently skipped when a company's AI activity originates from an eager R&D or IT team rather than from executive sponsorship. Isomorphic Labs offers a structural illustration of clean governance from the other side of the table: it "operates autonomously within Alphabet, with its own dedicated resources and exclusive focus on the application of AI to drug discovery," a deliberate separation from its parent company's broader AI business that gives its pharma partners a single, clearly scoped counterparty for governance and data-sharing purposes ([25] prnewswire.com).

Building Data Infrastructure and AI Readiness

No roadmap survives contact with fragmented data. Biotech data is generated across lab instruments, electronic lab notebooks, clinical systems, manufacturing execution systems, and external partners, and it is rarely stored in a form that is directly usable by a model. The Pistoia Alliance's 2025 Lab of the Future survey of more than 200 experts found that **data silos remain the top-cited challenge to making lab data usable, at 57% of respondents**, even though that figure has fallen nine percentage points since 2023 ([26]

pistoiaalliance.org). The same survey found cloud adoption in the lab climbing to **80% in 2025, up from 70% in 2023**, and electronic lab notebook (ELN) adoption rising sharply to **81% in 2025 from 66% in 2024** ([27] pistoiaalliance.org ([28] pistoiaalliance.org), a sign that the underlying plumbing is finally catching up to AI ambitions, though unevenly across organizations.

Compute infrastructure is the other half of readiness, and at the frontier it is now measured in GPU counts rather than server racks. Roche's March 2026 announcement of 2,176 additional Nvidia Blackwell GPUs brought its combined on-premises and cloud footprint past **3,500 Blackwell GPUs**, which Roche describes as the greatest announced GPU footprint available to any pharmaceutical company ([29] roche.com). Few biotechs will match that scale, and most should not try; renting cloud compute or partnering with a specialist AI vendor is the realistic path for the majority of mid-sized organizations, a point echoed by MIT's finding that **purchased AI tools and vendor partnerships succeed roughly 67% of the time, compared with far lower success rates for internally built systems** ([30] fortune.com). A practical readiness checklist a biotech should complete before its first roadmap milestone includes:

- **Data inventory and lineage**: cataloguing where scientific, clinical, and operational data lives, and who owns it.
- **Interoperability standards**: adopting common ontologies and schemas so that discovery, clinical, and manufacturing data can be queried together, an approach IntuitionLabs' own biotech data-readiness analysis frames as a prerequisite fix rather than an optional upgrade ([31] intuitionlabs.ai).
- **Compute strategy**: choosing between on-premises GPU clusters, cloud rental, or a hybrid model, sized to actual near-term use cases rather than aspirational ones.
- **Data governance and access controls**: establishing who can query, export, or fine-tune models on proprietary and patient data, ahead of any Good Clinical Practice (GCP) or Good Manufacturing Practice (GMP) validation requirement.
- **Vendor and partnership evaluation**: given the vendor-success advantage MIT documented, building a structured process to evaluate external AI partners rather than defaulting to in-house builds.

BCG's biopharma research puts the stakes of this checklist bluntly: **"without AI-ready data, biopharma companies will struggle to move beyond pilots,"** and leaders "invest in modern data platforms and stewardship to ensure everything from trial protocols to manufacturing records is AI-ready, so that initiatives can scale" ([32] bcg.com).

Selecting and Sequencing High-Value Use Cases

Use-case selection is where most roadmaps either gain or lose credibility with the board. McKinsey's deeper 2024 analysis, building on its initial economic estimate, identified more than 20 generative AI use cases with near-term potential across the life-sciences value chain, and its late-2024 survey found respondents concentrating investment in research functions (38%) more than any other area ([33] mckinsey.com). BCG's biopharma research quantifies the achievable range once a use case is scaled rather than piloted: leading firms have **accelerated clinical trials by up to 20%, improved manufacturing yields by more than 20%, and boosted the efficiency of sales representatives by 20% to 30%** ([34] bcg.com). BCG's related medtech research separately finds that AI-first medical-technology companies can increase revenues by 10% and productivity by 50% ([35] bcg.com).

Discovery remains the most crowded and most capital-intensive use case. Isomorphic Labs, the AlphaFold-descended digital-biology company spun out of Alphabet's DeepMind, received a **$45 million upfront payment with potential milestones up to $1.7 billion** in its first pharmaceutical partnership with Eli Lilly in January 2024 ([36] prnewswire.com), and later closed a **$2.1 billion Series B round** in May 2026 led by Thrive Capital, with

participants including Alphabet's GV, MGX, Temasek, and CapitalG, to scale its AI drug design engine ([37] intuitionlabs.ai). Alnylam Pharmaceuticals and Inceptive announced a foundation-model collaboration for RNA interference (RNAi) therapeutics valued at **up to $2 billion, with $30 million upfront**, explicitly designed to move "beyond rational drug design" ([38] biospace.com). What distinguishes Inceptive's platform from a narrower single-purpose model is that its "foundation model learns the patterns underlying biology and hence can adapt to diverse therapeutic modalities without retraining," which is the same generalization property that makes foundation models attractive for sequencing across multiple discovery use cases rather than building a bespoke model per program ([39] biospace.com). On clinical development, a peer-reviewed evaluation of an AI clinical trial matching system in Australian lung cancer patients, published via the National Institutes of Health's PubMed Central, found the system demonstrated **"96% accuracy in exclusion and 92% performance in assessing overall potential eligibility,"** while still requiring clinician oversight for nuanced cases ([40] pmc.ncbi.nlm.nih.gov), a meaningful improvement over manual screening for recruitment-constrained trials. On the commercial side, agentic AI is moving into customer relationship management (CRM): Veeva Systems, whose Vault CRM platform underpins commercial operations for many biopharma companies, projects roughly **15% to 20% industry efficiency improvements** from its embedded AI agents (intuitionlabs.ai).

The sequencing logic that emerges from this evidence favors starting with use cases that have a short feedback loop and a clear, measurable outcome metric, such as target identification triage or manufacturing yield optimization, before advancing to higher-risk, higher-regulatory-scrutiny applications such as autonomous clinical decision support. Owkin's 2026 expansion of its five-year "K Pro" agentic AI collaboration with Sanofi illustrates a staged approach: the companies' original 2021 collaboration was a **€90 million** partnership focused narrowly on oncology target identification and patient subgrouping, only later expanding into immunology drug positioning and now into autonomous multimodal biopharma agents ([42] owkin.com).

QIAGEN's May 2026 collaboration with Nvidia on graph-based AI for target and biomarker discovery offers a smaller-scale but instructive template for exactly this kind of staged sequencing. Rather than a broad launch, the companies explicitly designed **"initial pilot programs available to select pharmaceutical and biotechnology partners, with broader availability in future stages following validation"** ([43] biospace.com). QIAGEN's own justification for the collaboration rests on 25 years of curated biomedical knowledge underpinning more than 30,000 disease entries, on the theory that graph-based retrieval AI is only as trustworthy as the structured evidence it retrieves against, a validation-first philosophy directly transferable to any biotech evaluating a similar target-identification tool ([44] biospace.com).

Table 3 below situates six of the named partnerships and infrastructure investments discussed in this report side by side, illustrating the range of deal structures, from equity-linked collaborations to infrastructure buildouts to staged pilot programs, that biotechs can draw on when designing their own use-case sequencing.

Table 3: Selected Named AI Deals and Infrastructure Investments in Biotech and Pharma (2024-2026)

Companies	Deal Type	Disclosed Value or Scale	Primary Use Case
Novo Nordisk and OpenAI	Enterprise-wide integration	Full pipeline integration targeted by end of 2026 ([45] nexchron.com)	Discovery, manufacturing, and commercial forecasting
Roche and Nvidia	Infrastructure buildout	3,500+ Blackwell GPUs (see Case Studies, below)	Discovery, diagnostics, digital pathology
Alnylam and Inceptive	Strategic collaboration	Up to $2 billion, $30 million upfront ([38] biospace.com)	RNAi therapeutic design
Isomorphic Labs and Eli Lilly	Multi-target research collaboration	$45 million upfront, up to $1.7 billion in milestones ([36] prnewswire.com)	Small-molecule discovery

Companies	Deal Type	Disclosed Value or Scale	Primary Use Case
Isomorphic Labs Series B	Venture financing	$2.1 billion, led by Thrive Capital ([46] fiercebiotech.com)	Platform scale-up and pipeline expansion
Isomorphic Labs and Novartis	Multi-target research collaboration	$37.5 million upfront, up to $1.2 billion in milestones ([47] insideprecisionmedicine.com)	Small-molecule discovery across three targets
QIAGEN and Nvidia	Staged pilot collaboration	Initial pilots with select partners ([43] biospace.com)	Target and biomarker identification

Isomorphic's parallel deal with Novartis, announced the same day as its Lilly agreement, followed a similar milestone-heavy structure: **$37.5 million upfront with up to $1.2 billion in potential milestones** for small-molecule discovery against three undisclosed targets, described by Novartis president of biomedical research Fiona Marshall as harnessing "our companies' unique strengths, from AI and data science to medicinal chemistry and deep disease area expertise" ([48] insideprecisionmedicine.com).

The pattern across *Table 3* is that deal structure tends to track use-case risk: infrastructure and enterprise-wide integrations (Novo Nordisk, Roche) are undertaken by companies with the balance sheet to absorb a multi-year payback horizon, milestone-heavy collaborations (Alnylam, Isomorphic-Lilly) shift most of the financial risk onto the AI partner until value is proven, and staged pilots (QIAGEN-Nvidia) let both sides validate a narrower claim before wider rollout. A biotech without Roche's or Novo Nordisk's capital base should look to the latter two structures as the more realistic model for its own roadmap.

Governance, Risk Management, and Regulatory Compliance

Regulatory clarity has improved markedly since 2024 but remains incomplete, and a roadmap must build in compliance checkpoints rather than treating regulation as an afterthought. The U.S. Food and Drug Administration (FDA) issued its first draft guidance on AI in drug and biological product development in January 2025, establishing a **risk-based credibility assessment framework** for evaluating AI models against a defined "context of use" (COU) ([49] fda.gov). In January 2026, the FDA and the European Medicines Agency (EMA) jointly published **ten common principles for good AI practice across the medicines lifecycle**, covering research through post-market safety monitoring, described by European Commissioner Olivér Várhelyi as "a first step of a renewed EU-US cooperation" on novel medical technologies (ema.europa.eu (ema.europa.eu).

The EMA-FDA principles also point toward a broader legislative shift underway in Europe: the EMA notes that the new guiding principles build on the European Commission's Biotech Act proposal, which it says "holds great promise as a tool to accelerate the path from innovation to safe and effective medicines" and creates additional possibilities for testing innovative AI-driven methods in a controlled regulatory environment (ema.europa.eu). A biotech roadmap operating on a multi-year horizon should treat this as a live regulatory track to monitor alongside the AI Act itself, not a settled framework, particularly since the EMA notes that **"guideline development in the European Union (EU) is already underway, building on the EMA AI reflection paper published in 2024"** (ema.europa.eu).

The European Union's AI Act adds a second, more binding layer for companies operating in Europe. The Act (Regulation (EU) 2024/1689) entered into force on **1 August 2024** and is described by the European Commission as "the first-ever comprehensive legal framework on AI worldwide," using a risk-based structure with four tiers of risk, the strictest of which bans eight categories of prohibited practice outright (digital-strategy.ec.europa.eu (digital-strategy.ec.europa.eu). For biotech specifically, IntuitionLabs' analysis of the Act's pharma and medtech implications finds that AI systems embedded in or acting as a medical device under the EU Medical Device Regulation (MDR) or In Vitro Diagnostic Regulation (IVDR) are **automatically classified**

as high-risk, while AI used purely for internal pharmaceutical R&D or manufacturing generally falls outside that high-risk scope due to research exemptions in Articles 2(6) through 2(8) ([50] intuitionlabs.ai). The same analysis notes that full enforcement for AI embedded in regulated devices has effectively been extended to **August 2027** through a "Digital Omnibus" legislative proposal, giving device makers additional runway ([51] intuitionlabs.ai).

A governance program that satisfies both regulators and the board typically includes:

- **A model risk classification schema** mapping each AI use case to FDA credibility tiers or EU AI Act risk categories before development begins.
- **A documented context of use (COU) statement** for every AI model influencing a regulatory submission, per FDA's January 2025 draft guidance.
- **A cross-functional AI governance committee** spanning regulatory affairs, quality, IT, legal, and R&D leadership, mirroring the structures corporate-governance advisors recommend for boards generally ([52] trustible.ai).
- **Lifecycle monitoring commitments** for any AI model in production, since both the FDA and EMA frameworks extend oversight beyond initial validation into post-deployment monitoring.
- **Documented data provenance** for any training or fine-tuning data drawn from clinical or real-world sources, anticipating audit requests from regulators or acquirers.

Talent, Change Management, and Organizational Design

Even a well-governed, well-resourced roadmap stalls without people who can execute it, and this is where the evidence is most consistent across sources. A GlobalData survey covered by Pharmaceutical Executive found that a **shortage of specific skills and talent is the top hindrance to digital transformation** in pharma, a problem GlobalData's Urte Jakimaviciute describes as one companies "first experienced… shortly after the arrival of COVID-19," with "no quick solution" given how fast the underlying technology moves ([53] pharmexec.com). The Pistoia Alliance's 2025 survey corroborates this trend at the bench level, finding the share of respondents citing a lack of people as a barrier to AI adoption in the lab rose to **34%, up from 23% the prior year**, even as regulatory concern fell to just 9% of respondents from 23% ([54] pistoiaalliance.org), a genuine inversion of the industry's biggest worry over the space of a single year.

IntuitionLabs' 2025 skills-gap analysis, drawing on a GlobalData survey of 109 industry professionals, found that **49% of respondents identified a shortage of specific skills and talent as the top hindrance to digital transformation**, while a Pistoia Alliance survey separately found 44% of life-science R&D organizations citing a lack of skills as a major barrier to AI and machine learning adoption ([55] intuitionlabs.ai ([56] intuitionlabs.ai). On the response side, the same report finds that **reskilling existing employees delivered a 25% boost in retention and 15% efficiency gains**, at roughly half the cost of external hiring ([57] intuitionlabs.ai). MIT's research supports the same instinct from a different angle: because purchased AI tools and vendor partnerships succeed roughly twice as often as fully in-house builds, biotechs with thin internal AI talent are often better served hiring for AI literacy and integration skill (people who can adapt a vendor's tool to a regulated workflow) than for deep model-building expertise ([30] fortune.com). A practical talent plan should combine:

- **Reskilling tracks** for existing biologists, chemists, and clinical staff to gain applied AI literacy, prioritized over wholesale replacement.
- **Targeted external hiring** for a small number of "AI translator" roles who bridge biotech domain knowledge and data science.
- **Vendor and partnership scouting** as a formal function, not an ad hoc activity, given the higher success rate of bought-versus-built AI.

- **Change management ownership** assigned explicitly, since BCG's 10/20/70 finding places most of the adoption burden here rather than in the technology itself.

On the specific format of upskilling investment, the Pistoia Alliance's 2025 survey found clear preferences among life-science professionals themselves: **51% want best-practice guides, 45% want AI and machine learning courses, and 40% want structured skills training**, a signal that generic AI literacy content is less valued than material tailored to the regulatory and scientific context of the industry ([58] pistoiaalliance.org).

Implementation Considerations and Process Changes

Translating the five pillars above into an executable plan requires an operating rhythm, not a one-time announcement. Most organizations that scale successfully use a staged gate model: a discovery or diagnostic phase (typically 4 to 8 weeks) to inventory data and use cases; a pilot phase (typically one or two quarters) limited to one or two use cases with a pre-agreed success metric; a scale decision gate, where the executive sponsor and governance committee jointly decide whether to expand, redesign, or kill the pilot; and an enterprise rollout phase with lifecycle monitoring built in from day one. This mirrors what MIT's research calls the difference between organizations that "pick one pain point, execute well, and partner smartly" and those that spread thin across too many initiatives at once ([59] fortune.com).

Budget allocation deserves particular scrutiny, since MIT's research also found a **misalignment in resource allocation industry-wide: more than half of generative AI budgets go to sales and marketing tools, even though the biggest measured return on investment (ROI) came from back-office automation** such as eliminating outsourced business processes and cutting external agency costs ([60] fortune.com). A biotech roadmap should explicitly test that assumption against its own value chain rather than assume commercial or R&D use cases are automatically higher-value than operational ones.

Process changes should also anticipate the FDA's expectation, expressed in its January 2025 draft guidance and reinforced by the January 2026 joint FDA-EMA principles, that AI-influenced regulatory submissions carry documented evidence of model validation proportionate to risk ([61] fda.gov). Ark Biotech's use of the FDA credibility framework for a bioreactor simulation model, which IntuitionLabs' guidance analysis cites as an early real-world application, shows that even manufacturing-focused, non-clinical AI use cases now warrant the same documentation discipline as clinical ones once they influence a regulatory decision ([62] intuitionlabs.ai). Finally, organizations should build a formal review cadence, quarterly at minimum, that revisits the roadmap's use-case portfolio against new regulatory guidance, since both the FDA's guidance and the EU AI Act have moved substantially even within the 18 months preceding this report (digital-strategy.ec.europa.eu).

Documentation discipline extends to vendor governance as well. As adoption shifts toward agentic systems, IntuitionLabs' own review of Veeva's AI roadmap notes that the shift from a 2019-era embedded assistant to full autonomous agents "demand new data, governance, and change management processes," a caution that applies equally to any biotech evaluating a third-party agentic platform for clinical, regulatory, or commercial use rather than building one internally ([63] intuitionlabs.ai). A practical process checklist that follows from this evidence includes:

- **A written escalation path** for any AI output that a human reviewer disputes, defined before the system goes live, not improvised afterward.
- **Version-controlled model documentation**, recording every retraining or fine-tuning event against the FDA's lifecycle-monitoring expectation.

- **A named data steward** for each major data domain (discovery, clinical, manufacturing, commercial) accountable for data quality feeding AI systems.
- **A vendor exit plan** for every externally sourced AI tool, given how quickly the competitive landscape among AI vendors is still shifting.

Data Analysis and Evidence

The quantitative picture of AI adoption in biotech is one of accelerating investment paired with wide measurement uncertainty about ROI. *Table 1* below summarizes five independent market-size forecasts for the AI in drug discovery market, each using different base years, scope definitions, and methodologies, which is precisely why a roadmap should treat any single market number as a input to be triangulated rather than a fact to be quoted verbatim.

Table 1: AI in Drug Discovery Market Size Forecasts by Research Firm (as of July 2026)

Research Firm	Base Year Value	2026 Estimate	Long-Range Forecast	CAGR
Global Market Insights	$3.1 billion (2025)	$4 billion	$43.9 billion by 2035	30.5% ([64] gminsights.com)
Roots Analysis	$6.0 billion (2025)	$8.6 billion	$25.0 billion by 2035	12.6% ([65] rootsanalysis.com)
IMARC Group	$2.2 billion (2025)	n/a	$14.5 billion by 2034	22.70% ([3] imarcgroup.com)
Arizton	$1.71 billion (2024)	n/a	$8.52 billion by 2030	30.58% ([66] arizton.com)
Market Research Future	$1.172 billion (2025)	n/a	$11.82 billion by 2035	26.0% ([67] marketresearchfuture.com)

The five-fold spread in reported 2025 to 2026 market values, from roughly $1.2 billion to $8.6 billion, reflects differences in what each firm counts as “AI in drug discovery” (software only versus software plus services, narrow discovery tools versus broader R&D platforms) rather than a genuine disagreement about industry direction. All five forecasters agree on the trajectory: strong double-digit to high-double-digit growth through the early 2030s. For roadmap purposes, the practical takeaway is that the addressable market is large enough to justify sustained investment, but any specific market-size citation used in an internal business case should note its source and scope rather than be treated as a settled figure.

Adoption and execution data tell a more consistent story across sources, summarized in *Table 2*.

Table 2: Adoption, Governance, and Skills Metrics in Life Sciences AI (2024-2026)

Metric	Value	Source
Life sciences leaders who have experimented with generative AI	100% of surveyed leaders	McKinsey, late 2024 survey ([68] mckinsey.com)
Life sciences leaders who have scaled generative AI	32%	McKinsey, late 2024 ([69] mckinsey.com)
Life sciences leaders realizing consistent financial value	5%	McKinsey, late 2024 ([70] mckinsey.com)
Enterprise generative AI pilots reaching no measurable P&L impact (all industries)	~95%	MIT NANDA, 2025 ([71] fortune.com)

Metric	Value	Source
Life-science labs expecting to use AI within two years	77%	Pistoia Alliance, 2025 ([72] pistoiaalliance.org)
Labs citing skills shortage as a barrier to AI adoption	34% (up from 23% in 2024)	Pistoia Alliance, 2025 ([73] pistoiaalliance.org)
U.S. public companies disclosing board-level AI oversight	8% (245 of 3,048)	ISS-STOXX, January 2026 ([74] iss-stoxx.com)
U.S. public companies with a director holding specialized AI skills	16% (481 of 3,048)	ISS-STOXX, January 2026 ([75] iss-stoxx.com)
AI-first biopharma revenue uplift from full-scale adoption	5% to 15%	BCG, 2026 ([76] bcg.com)

Notably, this uneven track record has not dampened enthusiasm: McKinsey's same late-2024 survey found that **more than two-thirds of respondents planned to significantly increase their investment in generative AI** despite acknowledging how few organizations had converted early pilots into financial value ([77] mckinsey.com). Read together, these figures describe an industry where enthusiasm and early experimentation are nearly universal, but the ability to convert experimentation into durable financial value is concentrated in a small minority of organizations, roughly 5% by McKinsey's life-sciences-specific measure and MIT's cross-industry measure alike. That convergence across two independently run studies, one specific to life sciences and one spanning all sectors, is one of the more robust findings available to a roadmap author in 2026: whatever a biotech's specific circumstances, the base rate for AI pilots achieving a measurable, durable business outcome without a deliberate scaling strategy is low.

Deloitte's earlier but still-cited life-sciences AI survey adds useful texture on why value capture lags investment: it found the **top challenges impacting AI initiatives were difficulty identifying the business cases with the highest value (30%), data challenges (28%), and integrating AI into the organization (28%)** ([78] deloitte.com). Notably, all three top challenges are organizational and data-related rather than technical, reinforcing BCG's 10/20/70 framing that the algorithm itself is rarely the bottleneck. The same survey found that **43% of respondents had already used AI successfully to make processes more efficient**, the single most commonly realized benefit at that stage of industry adoption ([79] deloitte.com).

Deal activity, meanwhile, shows where capital is actually flowing even before ROI is proven at scale. Insilico Medicine's mid-2026 profit alert reported first-half 2026 revenue of **approximately $102.5 million to $106.5 million, a year-on-year increase of approximately 272.7% to 287.3%**, driven by a wave of out-licensing and collaboration deals with partners including Servier, Eli Lilly, and SK Biopharmaceuticals ([80] prnewswire.com). That is one of the clearest available signals that AI-native drug discovery, when paired with disciplined partnership strategy, can now generate real, growing, audited revenue rather than only speculative valuation.

Case Studies and Real-World Examples

Novo Nordisk and OpenAI: Full-Pipeline Integration

Novo Nordisk's 2026 partnership with OpenAI represents the most comprehensive publicly disclosed pharma-AI integration to date, spanning **dataset analysis, drug prospect identification, manufacturing process optimization, and sales and commercial forecasting**, with full integration targeted by the end of 2026 ([81] nexchron.com). The strategic driver is explicit: Novo Nordisk's glucagon-like peptide-1 (GLP-1) medicines, Ozempic and Wegovy, have driven extraordinary demand that has strained manufacturing capacity, while

competitors including Eli Lilly and Pfizer accelerate their own GLP-1 pipelines ([82] nexchron.com). Industry analysis of the deal notes that most prior pharmaceutical AI collaborations were narrow in scope, "a collaboration on protein modeling in a single therapeutic area, or a licensing arrangement for clinical trial optimization," which makes Novo Nordisk's commitment to "restructuring operating workflows across the entire value chain, including commercial functions" a departure in kind rather than merely in degree ([83] nexchron.com). This case demonstrates that the trigger for an enterprise-wide AI roadmap is often not AI capability itself but a business constraint, in this instance, capacity and competitive pressure, that AI is positioned to relieve across the full value chain rather than in a single department.

Roche and Nvidia: Building the AI Factory

Roche's expansion of its Nvidia-powered AI infrastructure illustrates the compute-first end of the roadmap spectrum. The company's **combined on-premises and cloud GPU infrastructure now exceeds 3,500 Blackwell GPUs**, built atop a strategic Nvidia collaboration that started in 2023 and now spans discovery, clinical trial design, digital pathology, manufacturing digital twins, and even guardrails for patient-facing digital health tools via Nvidia's NeMo Guardrails ([84] roche.com). Roche's own framing, that **"every day saved means a life-changing medicine or diagnostic reaches a patient sooner,"** according to Chief Digital and Technology Officer Wafaa Mamilli, ties the infrastructure investment directly to a clinical and commercial outcome rather than a technology-for-its-own-sake narrative ([85] roche.com). The same infrastructure also powers manufacturing digital twins, virtual replicas of production lines that let engineers optimize process and factory design before committing to physical changes, a lower-risk way to extract value from the same AI factory investment used for discovery. For most biotechs, Roche's specific GPU count is not a benchmark to match, but its governance discipline (naming a single accountable digital and technology executive for the entire program) is directly transferable.

Insilico Medicine: From AI Target Discovery to Phase III

Insilico Medicine's rentosertib program is among the clearest publicly documented examples of an AI-originated molecule reaching late-stage clinical development. The drug, an oral small-molecule inhibitor of TNIK, was identified through Insilico's Pharma.AI platform and reported as the top-ranked candidate in a *Nature Biotechnology* paper before phase IIa results were published in *Nature Medicine* and presented at the American Thoracic Society's 2025 conference; the company launched its **Phase III trial in July 2026 for idiopathic pulmonary fibrosis (IPF)**, enrolling 320 patients across roughly 50 sites, with a target New Drug Application (NDA) filing by March 2030 ([86] clinicalresearchnewsonline.com). Chief executive officer (CEO) Alex Zhavoronkov's own framing of the milestone, that AI typically compresses the path from target discovery to a developmental candidate to **"9-18 months"** but that later-stage clinical development still "moves with the speed of traffic," is one of the more candid public statements about where AI's speed advantage in biotech actually ends ([87] clinicalresearchnewsonline.com). This case is a useful corrective for any roadmap that implicitly assumes AI compresses the entire drug development timeline rather than specifically the early discovery phase.

The scale of Insilico's discovery engine is itself instructive for sequencing purposes: as of **June 30, 2026, the company's pipeline comprised 31 nominated preclinical candidate compounds (PCCs) across oncology, immunology, metabolism, and central nervous system disorders, with 13 having achieved investigational new drug (IND) clearance**, and 10 programs in active clinical development, split between four advanced independently and six run through partner collaborations ([88] insilico.com). One of the six PCCs nominated in the first half of 2026, an inhibitor developed by Insilico's United Arab Emirates (UAE) team, is described by the company as **"the first AI-driven innovative drug discovery milestone in the Gulf region,"** underscoring how

AI-native discovery platforms are also reshaping where, geographically, drug discovery capacity gets built ([89] insilico.com).

BenevolentAI and Recursion-Exscientia: Cautionary Tales in Scaling

Not every AI-native biotech's roadmap has produced its intended outcome, and a credible report should present these cases with equal weight. BenevolentAI, once valued at $2 billion after publicly claiming by 2018 to have built a **"bioscience machine brain, purpose-built to discover new medicines and cures for disease"** ([90] fiercebiotech.com), has undergone repeated rounds of layoffs since 2023, including a plan to cut **up to 180 people** after its lead pan-Trk inhibitor candidate failed a Phase 2a trial, followed by a further 30% staff reduction and closure of its U.S. office in 2024 ([91] fiercebiotech.com). In late 2024, the company announced a further "significant organizational restructuring" and a return to its "founding TechBio mission," shifting away from running its own late-phase clinical trials toward licensing standalone AI products to partners earlier in development ([92] fiercebiotech.com ([93] fiercebiotech.com).

Recursion Pharmaceuticals' 2024 merger with fellow AI-drug-discovery company Exscientia, intended to consolidate scale in the sector, was followed roughly six months later by a pipeline restructuring in which Recursion deprioritized three clinical-stage programs and paused another, narrowing its focus to six active oncology and rare-disease projects ([94] fiercebiotech.com). CEO Chris Gibson described the move as "a disciplined way of continuing on high-impact opportunities," rather than a retreat from AI-driven discovery itself ([95] fiercebiotech.com). The specific trigger for that retrenchment was clinical, not financial: a 62-patient Phase 2 trial of REC-994 in cerebral cavernous malformation showed only a **"trend towards reduced lesion volume and hemosiderin ring size"** at the higher dose, without meeting a bar the company judged sufficient to continue standalone ([96] fiercebiotech.com), a reminder that AI-originated candidates still face the same clinical proof bar as any other drug. Both cases illustrate the same underlying lesson for a roadmap author: an AI platform's technical credibility does not substitute for the ordinary biotech disciplines of capital efficiency, indication selection, and clinical execution, and a roadmap that treats "AI-native" as a strategy in itself, rather than a tool applied within a conventional biotech operating discipline, is vulnerable to the same setbacks as any other under-capitalized or overextended drug developer.

Owkin and Sanofi: Staged Expansion of an Agentic AI Partnership

Owkin's relationship with Sanofi shows what a multi-year, staged roadmap looks like from the vendor-partnership side. What began in 2021 as a **€90 million** collaboration focused narrowly on oncology target identification and patient subgrouping has, over five years, expanded first into immunology drug positioning and now into a new multi-year agreement to co-develop autonomous "biopharma agents" backed by a five-year license for Owkin's K Pro platform ([97] owkin.com). Sanofi's Chief Digital Officer Emmanuel Frenehard frames the expansion as part of a continuous investment strategy, stating the company aims "to empower our teams to operate with greater speed, depth, and confidence" through purpose-built agentic systems ([98] owkin.com). Owkin describes K Pro's design as combining "multimodal patient data with specialized biological, agentic AI systems to support each stage of the pharmaceutical value chain, from early discovery through clinical development," which is the same end-to-end ambition Novo Nordisk and Roche are pursuing at much larger scale ([99] owkin.com). This progression, narrow proof of value, then expansion by therapeutic area, then expansion by capability, is close to a textbook example of the sequencing logic this report recommends for biotechs building their own roadmaps with external AI partners rather than in-house platforms.

Implications and Future Directions

Several structural trends will shape how biotech AI roadmaps evolve over the next 12 to 24 months. First, regulatory convergence between the FDA and EMA, formalized in their January 2026 joint principles, suggests that companies operating across both U.S. and EU markets will increasingly be able to design a single global AI governance framework rather than maintaining fully separate regional programs, though the EU AI Act's additional high-risk obligations for MDR- and IVDR-linked AI mean full harmonization remains incomplete (ema.europa.eu). Companies should track the EU's "Digital Omnibus" process closely, since the intuitionlabs.ai analysis of the AI Act notes it has already shifted enforcement timelines once, from August 2026 to August 2027 for device-embedded AI ([100] intuitionlabs.ai). The Act's own risk taxonomy, spanning unacceptable, high, limited, and minimal risk tiers, means biotechs should classify each AI use case against that taxonomy explicitly rather than assume a blanket exemption for R&D activity (digital-strategy.ec.europa.eu).

Second, the consolidation-versus-partnership question in AI-native drug discovery remains unresolved. The Recursion-Exscientia merger sought scale through consolidation and required a subsequent pipeline retrenchment ([101] fiercebiotech.com), while Isomorphic Labs, Insilico Medicine, Owkin, and Inceptive have each pursued growth primarily through large-pharma partnerships and licensing deals rather than horizontal M&A. Harbour BioMed and BioMap's June 2026 launch of "MegaStream TechBio," a jointly created AI-native pipeline company combining proprietary datasets, purpose-built large models, and a dry-wet closed-loop discovery laboratory, suggests a third model is emerging: joint-venture creation of new AI-native entities rather than acquisition of existing ones ([102] prnewswire.com). The stated goal of that venture, to "systematically overcome the critical bottlenecks constraining next-generation innovative therapies," reads as a direct response to the very pilot-to-scale failure pattern this report has documented ([103] prnewswire.com). Biotechs building a roadmap should treat their choice among build, partner, and joint-venture as a first-order strategic decision, not a downstream implementation detail, particularly given MIT's finding that internally built AI systems succeed at roughly one-third the rate of purchased or partnered solutions ([30] fortune.com).

Third, the governance gap documented by ISS-STOXX is likely to close unevenly rather than uniformly. Health Care is already among the five sectors that, together with Industrials, Information Technology, Consumer Discretionary, and Financials, **"collectively represent nearly 75% of all disclosed instances of AI board oversight,"** meaning oversight is concentrating in a handful of sectors rather than diffusing broadly, but concentration data suggests early movers are pulling further ahead of laggards rather than the field converging ([104] iss-stoxx.com). Biotechs that establish board-level AI literacy and formal governance policies now are positioned to be counted among that leading group as institutional investors and proxy advisors increasingly scrutinize AI governance disclosures in the 2026 and 2027 proxy seasons.

Fourth, agentic AI, autonomous systems capable of executing multi-step tasks rather than simply generating text or predictions, is moving from pilot to production faster than prior generative AI waves. Veeva's own commercial rollout illustrates the pace: its first embedded AI assistant, Veeva Andi, launched in 2019 to give sales and marketing teams contextual "next best action" suggestions, was followed by CRM Bot, Voice, and MLR (medical, legal, regulatory) Bots in 2024, a formal Veeva AI initiative in 2025, and full AI Agents beginning deployment in December 2025, a five-year path from assistant to autonomous agent ([105] intuitionlabs.ai). Owkin's K Pro platform, Veeva's AI Agents in Vault CRM, and Novo Nordisk's OpenAI integration all point toward agentic systems handling increasingly consequential workflows, from drug positioning analysis to commercial forecasting, over the next two years. This raises the stakes for the governance and change-management pillars of any roadmap, since an autonomous agent making decisions inside a regulated workflow requires materially more validation and oversight than a model that simply drafts a document for human review.

Fifth, the market-sizing dispersion documented in *Table 1* is itself likely to narrow as research firms converge on standardized scope definitions. Mordor Intelligence's separate drug-discovery market analysis offers a data point on where AI specifically sits within the broader discovery market: **"AI and machine-learning platforms**

are projected to grow at 9.63% over 2026-2031," a narrower technology-segment view that helps explain why total-market forecasts vary so widely depending on whether "AI in drug discovery" is scoped as its own market or as one growing slice of the broader drug discovery market ([106] mordorintelligence.com). Finally, the persistent skills gap suggests that talent strategy, not compute or algorithms, will remain the binding constraint on how quickly individual biotechs can move from the 32% who have taken steps to scale (per McKinsey) toward the 5% who have realized durable financial value. Organizations that invest early and specifically in reskilling, rather than assuming external hiring alone will close the gap, are likely to compound an advantage as the AI talent market remains tight across the broader technology sector, not just life sciences.

Frequently Asked Questions (FAQs)

What is an AI roadmap for a biotech company? It is a sequenced, governed plan for moving from isolated AI experimentation to durable deployment across R&D, clinical, manufacturing, and commercial functions, typically anchored by executive sponsorship, a data-readiness assessment, a prioritized use-case portfolio, a governance and regulatory framework, and a talent and change-management plan.

How do you build a biotech AI strategy from scratch? Start with a data and use-case inventory, secure explicit executive and board sponsorship (given that only 8% of large U.S. companies currently disclose board-level AI oversight ([74] iss-stoxx.com)), select one or two pilots with clear success metrics, and design the governance and compliance framework before scaling, rather than after. This sequencing mirrors Deloitte's own conclusion that companies should "build a vision and strategy to leverage AI, then put in place the building blocks needed to scale its use," including the right information technology (IT) infrastructure, talent, and ecosystem partnerships, in that order ([107] deloitte.com).

What are the most common biopharma AI use cases? McKinsey identifies more than 20 near-term use cases spanning target identification, generative molecule design, clinical trial optimization and patient matching, manufacturing process control, medical writing, and commercial forecasting, with research functions currently receiving the largest share of investment ([33] mckinsey.com).

How should a biotech present its AI strategy to the board? Board decks should quantify the data-readiness gap, name an accountable executive sponsor, present a small prioritized use-case portfolio with expected value ranges grounded in independent benchmarks such as BCG's 5% to 15% AI-first revenue uplift figure, and include an explicit governance and regulatory compliance plan rather than treating AI oversight as implicit within existing IT governance.

How is AI governed in life sciences, and what regulatory guidance applies? The FDA's January 2025 draft guidance establishes a risk-based credibility framework tied to a model's context of use; the January 2026 FDA-EMA joint principles harmonize expectations across the medicines lifecycle; and the EU AI Act imposes additional binding obligations on AI embedded in or acting as a regulated medical device, with most high-risk obligations enforceable from August 2026, extended to August 2027 for device-embedded AI under the pending Digital Omnibus proposal.

Why do most AI pilots in pharma and biotech fail to scale? MIT's 2025 research attributes roughly 95% of enterprise generative AI pilot failures, across all industries, to a "learning gap" in how tools are integrated into actual workflows rather than to model quality, and finds that purchased or partnered AI solutions succeed roughly 67% of the time compared with substantially lower success for internally built systems ([108] fortune.com).

How does a life-sciences AI consulting engagement typically start? Most engagements begin with a data and process diagnostic phase, benchmarking a company's data infrastructure, governance maturity, and use-case portfolio against industry evidence of what has and has not scaled elsewhere, before recommending a

prioritized, gated roadmap rather than a single large-scale AI deployment. As an adjacent advisor rather than a software vendor, IntuitionLabs approaches this diagnostic the same way its own published analyses treat the industry: grounding every recommendation in verifiable regulatory guidance, market data, and named case outcomes rather than platform-specific sales claims ([109] intuitionlabs.ai).

What data infrastructure does a biotech need before starting an AI roadmap? At minimum, a unified data inventory with clear lineage, common ontologies or schemas across research and clinical systems, a defined compute strategy (on-premises, cloud, or hybrid), and documented access controls, since fragmented and siloed data remains the top-cited obstacle to scaling AI in the lab, cited by 57% of life-science professionals surveyed by the Pistoia Alliance in 2025 ([26] pistoiaalliance.org).

Conclusion

Building an AI roadmap for a biotech company in 2026 requires resisting two opposite temptations: treating AI as a single transformative initiative to be announced and funded all at once, or treating it as a collection of disconnected departmental pilots with no shared governance or sequencing logic. The evidence assembled in this report, from McKinsey's finding that only 5% of life-sciences organizations have converted generative AI experimentation into consistent financial value, to MIT's parallel finding of a 95% cross-industry pilot failure rate, to BCG's 10/20/70 framing of where the real adoption challenge lies, points toward a narrower and more disciplined path: align the board and executive team around a small number of prioritized use cases, fix the underlying data and compute infrastructure before scaling, build regulatory and governance guardrails proportionate to each use case's risk rather than uniformly, and invest specifically in reskilling and change management rather than assuming technology alone will carry the transformation.

The named cases in this report, Roche's infrastructure-led approach, Novo Nordisk's full-pipeline OpenAI integration, Insilico Medicine's discovery-to-Phase-III trajectory, Owkin and Sanofi's staged multi-year partnership, and the more cautionary trajectories of BenevolentAI and the Recursion-Exscientia merger, together demonstrate that no single operating model guarantees success, but that discipline in sequencing, capital allocation, and governance consistently distinguishes the organizations converting AI investment into durable outcomes from those absorbing repeated setbacks. For a biotech leadership team building its first formal AI roadmap, or revising one built before the FDA's January 2025 guidance and the EU AI Act's phased enforcement became clear, the sequencing logic in this report, strategy and governance first, data infrastructure second, prioritized use cases third, regulatory alignment fourth, and talent and change management as the continuous thread running through all four, offers a evidence-grounded starting structure rather than a generic technology adoption checklist.

External Sources

[1] https://www.gminsights.com/industry-analysis/ai-in-drug-discovery-market#:~:The%2...

[2] https://www.rootsanalysis.com/reports/ai-based-drug-discovery-market.html#:~:The%2...

[3] https://www.imarcgroup.com/ai-in-drug-discovery-market#:~:The%2...

[4] https://www.deloitte.com/us/en/industries/life-sciences-health-care/articles/value-of-genai-in-pharma.html#:~:The%2...

[5] https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry#:~:All%2...

[6] https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/#:~:But%2...

[7] https://www.bcg.com/publications/2026/scaling-ai-in-the-biopharmaceutical-industry#:~:Leade...

[8] https://www.bcg.com/publications/2026/scaling-ai-in-the-biopharmaceutical-industry#:~:a%205...

[9] https://www.iss-stoxx.com/insights/articles/mind-the-governance-gap-the-state-of-board-oversight-and-ai-policy-in-us-companies/#:~:Out%2...

[10] https://pistoiaalliance.org/news/survey-ai-adoption-life-sciences-labs-skills-gap/#:~:Howev...

[11] https://www.roche.com/media/releases/med-cor-2026-03-16#:~:With%...

[12] https://www.clinicalresearchnewsonline.com/news/2026/07/08/insilico-medicine-launches-phase-iii-trial-for-ai-developed-drug#:~:Insil...

[13] https://www.fiercebiotech.com/biotech/benevolentai-pivots-back-techbio-roots-causing-more-layoffs#:~:Benev...

[14] https://intuitionlabs.ai/articles/pharma-ai-skills-gap#:~:one%2...

[15] https://nexchron.com/health/novo-nordisk-openai-drug-pipeline-integration-2026#:~:Novo%...

[16] https://www.roche.com/media/releases/med-cor-2026-03-16#:~:In%20...

[17] https://www.deloitte.com/us/en/insights/industry/life-sciences/ai-and-pharma.html#:~:More%...

[18] https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry#:~:resea...

[19] https://www.prnewswire.com/news-releases/isomorphic-labs-announces-strategic-multi-target-research-collaboration-with-lilly-302027392.html#:~:That%...

[20] https://www.mordorintelligence.com/industry-reports/drug-discovery-market#:~:High-...

[21] https://www.bcg.com/publications/2026/scaling-ai-in-the-biopharmaceutical-industry#:~:They%...

[22] https://www.bcg.com/publications/2026/scaling-ai-in-the-biopharmaceutical-industry#:~:These...

[23] https://www.iss-stoxx.com/insights/articles/mind-the-governance-gap-the-state-of-board-oversight-and-ai-policy-in-us-companies/#:~:This%...

[24] https://www.deloitte.com/us/en/insights/industry/life-sciences/ai-and-pharma.html#:~:As%20...

[25] https://www.prnewswire.com/news-releases/isomorphic-labs-announces-strategic-multi-target-research-collaboration-with-lilly-302027392.html#:~:Isomo...

[26] https://pistoiaalliance.org/news/survey-ai-adoption-life-sciences-labs-skills-gap/#:~:Data%...

[27] https://pistoiaalliance.org/news/survey-ai-adoption-life-sciences-labs-skills-gap/#:~:Cloud...

[28] https://pistoiaalliance.org/news/survey-ai-adoption-life-sciences-labs-skills-gap/#:~:ELNs%...

[29] https://www.roche.com/media/releases/med-cor-2026-03-16#:~:this%...

[30] https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/#:~:Purch...

[31] https://intuitionlabs.ai/articles/ai-data-infrastructure-biotech#:~:Biote...

[32] https://www.bcg.com/publications/2026/scaling-ai-in-the-biopharmaceutical-industry#:~:Witho...

[33] https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry#:~:Six%2...

[34] https://www.bcg.com/publications/2026/scaling-ai-in-the-biopharmaceutical-industry#:~:accel...

[35] https://www.bcg.com/publications/2025/the-ai-first-medical-technology-company?recommendedArticles=true#:~:Mede...

[36] https://www.prnewswire.com/news-releases/isomorphic-labs-announces-strategic-multi-target-research-collaboration-with-lilly-302027392.html#:~:Isomo...

[37] https://intuitionlabs.ai/articles/isomorphic-labs-series-b-ai-drug-discovery#:~:Isomo...

[38] https://www.biospace.com/press-releases/alnylam-and-inceptive-form-strategic-ai-collaboration-to-accelerate-the-discovery-of-rnai-therapeutics#:~:Colla...

[39] https://www.biospace.com/press-releases/alnylam-and-inceptive-form-strategic-ai-collaboration-to-accelerate-the-discovery-of-rnai-therapeutics#:~:Incep...

[40] https://pmc.ncbi.nlm.nih.gov/articles/PMC7382632/#:~:This%...

[41] https://intuitionlabs.ai/articles/veeva-ai-roadmap-crm-bot-agents-2026#:~:these...

[42] https://www.owkin.com/newsfeed/owkin-to-build-ai-agents-as-part-of-a-multi-year-k-pro-collaboration-with-sanofi#:~:Owkin...

[43] https://www.biospace.com/press-releases/qiagen-to-advance-ai-driven-drug-discovery-with-graph-based-ai-and-curated-bioinformatics-knowledge-with-nvidia#:~:Initi...

[44] https://www.biospace.com/press-releases/qiagen-to-advance-ai-driven-drug-discovery-with-graph-based-ai-and-curated-bioinformatics-knowledge-with-nvidia#:~:By%20...

[45] https://nexchron.com/health/novo-nordisk-openai-drug-pipeline-integration-2026#:~:Novo%...

[46] https://www.fiercebiotech.com/biotech/alphabets-ai-biotech-isomorphic-labs-bags-21b-series-b-fuel-next-gen-drug-design-model#:~:Thriv...

[47] https://www.insideprecisionmedicine.com/topics/precision-medicine/deepminds-isomorphic-inks-3b-worth-of-deals-with-lilly-and-novartis-in-one-day/#:~:the%2...

[48] https://www.insideprecisionmedicine.com/topics/precision-medicine/deepminds-isomorphic-inks-3b-worth-of-deals-with-lilly-and-novartis-in-one-day/#:~:This%...

[49] https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological#:~:this%...

[50] https://intuitionlabs.ai/articles/eu-ai-act-pharma-medical-device-compliance#:~:AI%20...

[51] https://intuitionlabs.ai/articles/eu-ai-act-pharma-medical-device-compliance#:~:a%20%...

[52] https://trustible.ai/post/how-to-establish-an-effective-ai-governance-committee-in-2026/#:~:The%2...

[53] https://www.pharmexec.com/view/survey-struggling-digital-transformation#:~:Short...

[54] https://pistoiaalliance.org/news/survey-ai-adoption-life-sciences-labs-skills-gap/#:~:Regul...

[55] https://intuitionlabs.ai/articles/pharma-ai-skills-gap#:~:in%20...

[56] https://intuitionlabs.ai/articles/pharma-ai-skills-gap#:~:a%20P...

[57] https://intuitionlabs.ai/articles/pharma-ai-skills-gap#:~:one%2...

[58] https://pistoiaalliance.org/news/survey-ai-adoption-life-sciences-labs-skills-gap/#:~:Risin...

[59] https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/#:~:%E2%8...

[60] https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/#:~:The%2...

[61] https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological#:~:this%...

[62] https://intuitionlabs.ai/articles/fda-ai-drug-development-guidance#:~:an%20...

[63] https://intuitionlabs.ai/articles/veeva-ai-roadmap-crm-bot-agents-2026#:~:but%2...

[64] https://www.gminsights.com/industry-analysis/ai-in-drug-discovery-market#:~:The%2...

[65] https://www.rootsanalysis.com/reports/ai-based-drug-discovery-market.html#:~:The%2...

[66] https://www.arizton.com/market-reports/ai-in-drug-discovery-market#:~:THE%2...

[67] https://www.marketresearchfuture.com/reports/ai-drug-discovery-market-9393#:~:The%2...

[68] https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry#:~:All%2...

[69] https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry#:~:32%20...

[70] https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry#:~:only%...

[71] https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/#:~:the%2...

[72] https://pistoiaalliance.org/news/survey-ai-adoption-life-sciences-labs-skills-gap/#:~:the%2...

[73] https://pistoiaalliance.org/news/survey-ai-adoption-life-sciences-labs-skills-gap/#:~:skill...

[74] https://www.iss-stoxx.com/insights/articles/mind-the-governance-gap-the-state-of-board-oversight-and-ai-policy-in-us-companies/#:~:only%...

[75] https://www.iss-stoxx.com/insights/articles/mind-the-governance-gap-the-state-of-board-oversight-and-ai-policy-in-us-companies/#:~:only%...

[76] https://www.bcg.com/publications/2026/scaling-ai-in-the-biopharmaceutical-industry#:~:a%205...

[77] https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry#:~:Nonet...

[78] https://www.deloitte.com/us/en/insights/industry/life-sciences/ai-and-pharma.html#:~:Top%2...

[79] https://www.deloitte.com/us/en/insights/industry/life-sciences/ai-and-pharma.html#:~:Most%...

[80] https://www.prnewswire.com/news-releases/insilico-medicine-releases-positive-profit-alert-for-the-first-half-of-2026-302821358.html#:~:Insil...

[81] https://nexchron.com/health/novo-nordisk-openai-drug-pipeline-integration-2026#:~:Accor...

[82] https://nexchron.com/health/novo-nordisk-openai-drug-pipeline-integration-2026#:~:The%2...

[83] https://nexchron.com/health/novo-nordisk-openai-drug-pipeline-integration-2026#:~:What%...

[84] https://www.roche.com/media/releases/med-cor-2026-03-16#:~:With%...

[85] https://www.roche.com/media/releases/med-cor-2026-03-16#:~:In%20...

[86] https://www.clinicalresearchnewsonline.com/news/2026/07/08/insilico-medicine-launches-phase-iii-trial-for-ai-developed-drug#:~:Insil...

[87] https://www.clinicalresearchnewsonline.com/news/2026/07/08/insilico-medicine-launches-phase-iii-trial-for-ai-developed-drug#:~:Usual...

[88] https://insilico.com/news/6f8yhr6sl1-insilico-medicine-releases-positive-prof?amp=true#:~:As%20...

[89] https://insilico.com/news/6f8yhr6sl1-insilico-medicine-releases-positive-prof?amp=true#:~:repre...

[90] https://www.fiercebiotech.com/biotech/benevolentai-pivots-back-techbio-roots-causing-more-layoffs#:~:By%20...

[91] https://www.fiercebiotech.com/biotech/benevolentai-pivots-back-techbio-roots-causing-more-layoffs#:~:Benev...

[92] https://www.fiercebiotech.com/biotech/benevolentai-pivots-back-techbio-roots-causing-more-layoffs#:~:Benev...

[93] https://www.fiercebiotech.com/biotech/benevolentai-pivots-back-techbio-roots-causing-more-layoffs#:~:Going...

[94] https://www.fiercebiotech.com/biotech/several-months-after-exscientia-merge-ai-outfit-recursion-reworks-pipeline#:~:All%2...

[95] https://www.fiercebiotech.com/biotech/several-months-after-exscientia-merge-ai-outfit-recursion-reworks-pipeline#:~:Rare%...

[96] https://www.fiercebiotech.com/biotech/several-months-after-exscientia-merge-ai-outfit-recursion-reworks-pipeline#:~:MRI%2...

[97] https://www.owkin.com/newsfeed/owkin-to-build-ai-agents-as-part-of-a-multi-year-k-pro-collaboration-with-sanofi#:~:Owkin...

[98] https://www.owkin.com/newsfeed/owkin-to-build-ai-agents-as-part-of-a-multi-year-k-pro-collaboration-with-sanofi#:~:By%20...

[99] https://www.owkin.com/newsfeed/owkin-to-build-ai-agents-as-part-of-a-multi-year-k-pro-collaboration-with-sanofi#:~:K%20P...

[100] https://intuitionlabs.ai/articles/eu-ai-act-pharma-medical-device-compliance#:~:AI%20...

[101] https://www.fiercebiotech.com/biotech/several-months-after-exscientia-merge-ai-outfit-recursion-reworks-pipeline#:~:The%2...

[102] https://www.prnewswire.com/news-releases/harbour-biomed-and-biomap-jointly-initiate-megastream-techbio-a-leading-global-drug-development-platform-joins-forces-with-premier-life-science-foundation-models-to-set-new-benchmarks-for-ai-driven-complex-biologics-rd-302800277.html#:~:the%2...

[103] https://www.prnewswire.com/news-releases/harbour-biomed-and-biomap-jointly-initiate-megastream-techbio-a-leading-global-drug-development-platform-joins-forces-with-premier-life-science-foundation-models-to-set-new-benchmarks-for-ai-driven-complex-biologics-rd-302800277.html#:~:The%2...

[104] https://www.iss-stoxx.com/insights/articles/mind-the-governance-gap-the-state-of-board-oversight-and-ai-policy-in-us-companies/#:~:The%2...

[105] https://intuitionlabs.ai/articles/veeva-ai-roadmap-crm-bot-agents-2026#:~:The%2...

[106] https://www.mordorintelligence.com/industry-reports/drug-discovery-market#:~:AI%20...

[107] https://www.deloitte.com/us/en/insights/industry/life-sciences/ai-and-pharma.html#:~:compa...

[108] https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/#:~:The%2...

[109] https://intuitionlabs.ai/articles/pharma-ai-skills-gap#:~:This%...

IntuitionLabs - Industry Leadership & Services

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

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

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

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

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

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

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

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

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

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

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

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

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

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

DISCLAIMER

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

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

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

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

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

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