

Build vs Buy AI in Pharma: R&D and Commercial Guide

2/5/2026 • 30 min read

- ai in pharma
- build vs buy strategy
- pharmaceutical r&d
- drug discovery ai
- clinical trials
- commercial pharma
- artificial intelligence
- cost benefit analysis
- digital transformation
- machine learning



Executive Summary

The pharmaceutical industry stands at a crossroads in its adoption of artificial intelligence (AI). Rapid advances in AI promise to accelerate drug discovery, optimize clinical trials, personalize medicine, and transform commercial operations, but they also raise the pivotal **build-vs-buy** decision for R&D and commercial teams. Building AI in-house offers maximal customization and control over proprietary data and IP, whereas buying or partnering with vendors provides speed, lower upfront cost, and access to specialized expertise. Our analysis finds that **hybrid strategies** are increasingly favored: **start by buying** to achieve quick wins and validate use cases, then **build custom systems** where in-house differentiation and long-term value justify the investment ⁽¹⁾ intuitionlabs.ai ⁽²⁾ www.zartis.com. This report provides a comprehensive framework to guide pharmaceutical leaders through that decision, covering: (a) the state of AI in pharma R&D and commercialization; (b) strategic considerations such as cost, time-to-value, risk, and competitive differentiation; (c) vendor evaluation and total cost of ownership; (d) data, regulatory, and organizational factors; and (e) case studies from industry. We show that while **speed-to-impact** often points to off-the-shelf solutions (e.g. purchasing AI-powered drug-screening platforms or clinical trial management software), in-house AI can yield **proprietary advantage** for core pipelines. Ultimately, the optimal path depends on aligning AI adoption with corporate strategy: if an AI capability is **core to your value proposition**, building it in-house (or through a close partnership) may pay off; if it is **supportive or commodity**, buying is usually wiser ⁽³⁾ www.linkedin.com ⁽⁴⁾ intuitionlabs.ai. The report concludes with practical guidelines, tables, and examples to equip R&D and commercial teams to decide wisely.

Introduction

The pharmaceutical industry is characterized by enormously complex, time-consuming, and capital-intensive R&D and commercialization processes. Bringing a new drug to market typically takes over a decade and often costs on the order of **billion-dollar** investments ⁽⁵⁾ apnews.com. Outcomes are highly uncertain: out of thousands of early discovery projects, only a few yield an approved therapy. In this context, AI and data-driven tools hold great promise to improve efficiency and productivity across the R&D-to-commercial continuum. Major players are therefore investing heavily: for example, **AstraZeneca's** Chief Data Scientist notes that *"data science and AI are transforming R&D, helping us turn science into medicine more quickly and with a higher probability of success"* ⁽⁶⁾ moneyweek.com. Likewise, **GSK** announced a ~\$30 billion U.S. investment that explicitly includes deploying AI in manufacturing and R&D workflows ⁽⁷⁾ moneyweek.com.

Despite this surge of interest, the pharma sector faces a central strategic dilemma: **should AI capabilities be built in-house or bought from external providers?** Many large pharma companies lack deep software development expertise and face chronic talent shortages in AI and data science. Outsourcing or buying AI solutions can give immediate access to cutting-edge technology. On the other hand, in-house development can protect proprietary data and tailor solutions precisely to the company's pipelines, potentially yielding a lasting competitive edge. This build-vs-buy choice is non-trivial. If done poorly, companies can waste millions on failed custom projects or slide into *"pilot purgatory"*, where dozens of small AI pilots never scale to meaningful impact ⁽⁸⁾ intuitionlabs.ai ⁽⁹⁾ www.pharmexec.com. This report provides an in-depth analysis of that decision framework, with a balanced view for both **R&D teams** (drug discovery, clinical trials, manufacturing, etc.) and **commercial teams** (sales, marketing, patient engagement, supply chain). We integrate data, expert opinions, and case studies to help pharma leaders decide when to prototype or buy, and when—and how—to invest in custom AI capabilities.

AI Adoption in Pharma: Background and Trends

Pharma and biotech companies have historically lagged behind other industries in digital innovation, partly due to heavy regulation and legacy processes. Over the past decade, however, the explosion of biomedical and patient data and the

maturity of machine learning models have begun transforming pharmaceutical science. In drug discovery, biotech startups and academic consortia have used AI to analyze genetic, proteomic, and chemical datasets for target identification or novel molecule design. For example, **AlphaFold2** – a deep learning system by Google DeepMind – can predict a protein’s 3D structure from its amino acid sequence, a breakthrough that helps researchers infer drug targets ([10] moneyweek.com). In 2024, Insitro (an AI-focused biotech) signed collaboration deals with Eli Lilly and Bristol Myers Squibb, using ML to analyze large “**chemical and biological markers**” datasets in search of new metabolic and neurological disease therapies ([11] apnews.com). In clinical development, AI startups like Formation Bio claim to speed up trials by “as much as 50%” through automating patient screening, regulatory filings, and matching therapies to diseases ([12] time.com).

On the commercial side, digital transformation accelerated during and after COVID-19. Remote sales and telehealth became mainstream, and pharma organizations sought AI to personalize customer interactions. A recent survey found about **75%** of life sciences executives reported implementing some form of AI in the past 2 years, with another 86% planning rollouts within two years ([13] www.axios.com). For example, CRM platforms tailored for life sciences (like Veeva or Workbooks) are embedding AI features. Workbooks, a UK-based CRM used by sales teams, recently integrated generative-AI “Sales Coach” and transcription tools to automate repetitive tasks and improve targeting ([14] www.techradar.com). Meanwhile, generative AI has started influencing content creation: pharma marketers experiment with AI for personalized patient education and HCP communications, although concrete data on ROI is still maturing.

At the ecosystem level, an interesting trend noted by analysts is that **large incumbents and small biotech startups play complementary roles** in the AI transformation. Causeway Capital reports that small, nimble teams wielding AI can compress discovery time, but big pharma still retains scale advantages in later stages ([15] moneyweek.com). In practice, large companies often partner with or acquire AI startups to tap their expertise, while also internalizing successful approaches. The **key question** for any R&D or commercial leader is identifying that inflection point: “Will an AI initiative give us a unique advantage worth building, or is it better to leverage an existing solution?” ([3] www.linkedin.com).

The Build-vs-Buy Dilemma in Pharma AI

Pharma companies deciding between building or buying AI solutions must weigh multiple factors that affect both short-term success and long-term strategy. Table 1 below summarizes key differences in cost, timeline, and risk. Building custom AI systems in-house typically incurs **high up-front investment** (on the order of hundreds of thousands to millions of dollars) and substantial ongoing maintenance costs ([16] xenoss.io). Hiring AI talent alone can mean annual salaries in the high six figures per expert, plus GPU infrastructure costs (on the order of \$0.2–2M/year) ([17] intuitionlabs.ai). In contrast, buying a commercial AI platform often requires a relatively modest license fee (tens of thousands per year) and a smaller annual subscription, shifting many R&D costs to the vendor. However, those vendor solutions offer **limited customization**: you must adapt your workflows to the software rather than the other way around.

Factor	Build (In-House)	Buy (Commercial/Platform)	Hybrid/Partnership
Initial Investment	High (~\$500K–\$2M+) per project ([16] xenoss.io)	Low (~\$50K–\$200K) for licensing/common tools ([18] xenoss.io)	Medium (~\$100K–\$500K) for collaboration ([16] xenoss.io)
Ongoing Cost (% Ann.)	High (30–40% of development cost) ([16] xenoss.io)	Low-Medium (10–20%) ([18] xenoss.io)	Medium (15–25%) ([19] xenoss.io)
Time to Value	Long (typically 12–24 months or more) ([16] xenoss.io)	Short (3–6 months to deployment) ([18] xenoss.io)	Intermediate (6–12 months) ([16] xenoss.io)
Control / Flexibility	Maximum (full customization and IP control) ([20] intuitionlabs.ai)	Limited (constrained to vendor’s features)	Shared (joint solution) ([20] intuitionlabs.ai)
Risk Profile	High (technical and implementation risk)	Lower (tech risk is reduced, but risk of vendor lock-in) ([21] intuitionlabs.ai)	Medium (coordination/startegic risk)

Table 1. Comparison of key factors when building AI in-house vs buying external solutions. (Cost and risk ranges adapted from industry analyses (^[16] xenoss.io) (^[20] intuitionlabs.ai)).

From the above, certain patterns emerge:

- **Speed and Time-to-Value:** Commercial AI products (or cloud services) can often be deployed in weeks or a few months, giving quick proof-of-concept. This rapid turnaround is valuable for demonstrating ROI and building momentum. (^[9] intuitionlabs.ai) (^[18] xenoss.io) Conversely, developing an in-house system can take a year or more, during which strategic priorities may shift or funding may run out, leading to “pilot purgatory” (^[9] intuitionlabs.ai).
- **Cost and ROI:** Building in-house entails significant up-front costs: not only salaries for data scientists and engineers (on the order of \$200K–\$500K annually per senior specialist (^[22] xenoss.io)) but also expenses for computing, data curation, and compliance. Industry estimates suggest an in-house AI project might require \$0.5–\$2M initially plus ~30–40% of that each year, whereas using a commercial AI platform could cost on the order of \$50K–\$200K upfront and 10–20% annually (^[4] intuitionlabs.ai). Importantly, commercial vendors absorb many hidden costs (e.g. ongoing model maintenance, software updates, regulatory documentation) that would otherwise fall on the company.
- **Strategic Control and IP:** If the AI capability is core to a firm's competitive advantage (e.g. a novel drug-design algorithm or patient stratification engine that could be patented), building in-house preserves control and ownership. As one industry expert noted, “*full ownership of the system itself means you can modify it as you see fit, patent innovative solutions to protect your competitive advantage, and even market your system as a proprietary solution*” (^[23] www.linkedin.com). Buying off-the-shelf relinquishes much of that control. Even if a vendor prohibits data misuse, using their model means you share insights (implicitly) with the provider's systems, which may accumulate knowledge over time (^[23] www.linkedin.com). Thus, if data ownership or regulatory compliance absolutes are paramount, building (or tightly partnering) may be necessary.
- **Expertise and Maintenance:** Building sophisticated AI systems requires scarce talent. If an organization already has experienced AI/ML teams (e.g. a tech-centric pharma or big biotech), developing in-house may be feasible. Many pharma, however, rely on partnerships. Vendor solutions often come with built-in maintenance (automated updates, monitoring dashboards, support). For example, a SaaS machine learning tool might automatically retrain models on new data, whereas an in-house model would require building a full MLOps pipeline (^[24] intuitionlabs.ai). The ongoing operational burden can be major: after about a year custom AI costs often intensify as hidden elements like CI/CD, security, and compliance become necessary (^[24] intuitionlabs.ai).
- **Risk and Compliance:** Regulatory risk is high in pharma. If an AI solution handles patient-level or clinical trial data, both building and buying require strict validation. Vendors may come with pre-built compliance (HIPAA, GxP logging, audit trails), but failing to scrutinize a vendor's processes can create liability. Conversely, building in-house means bearing the entire burden of validation and audits. In high-stakes domains, many companies prefer tested domain-specific vendors with proven compliance track record.

In summary, **buying is often the pragmatic choice for non-differentiating needs (e.g., broad data analytics, commercial tools)**, while **building may be justified for strategically unique capabilities**. In practice, firms frequently blend approaches: validating concepts with vendor tools and later investing in custom development where it yields sufficient returns (^[1] intuitionlabs.ai) (^[2] www.zartis.com). The sections below unpack these considerations in greater depth for R&D and commercial teams.

Key Considerations for Build vs Buy

When evaluating any specific AI initiative, pharma decision-makers should systematically weigh a set of criteria:

- **Strategic Importance and Core Contribution:** Is the AI capability a *core* feature of a future product or service? If so, in-house development may be critical. Dr. Andrée Bates (former Roche and AstraZeneca data scientist) suggests asking: “*Will the new system be a core feature or critical function of your offering? If so, it's likely better to build,*” whereas for supporting functions (e.g. marketing analytics), buying suffices (^[3] www.linkedin.com). In R&D, predicting therapeutic outcomes or discovering fundamentally new biology might be core, while routine data reporting is not. In commercial, predictive ad targeting may be strategic, but many aspects (like SMS reminders) can use third-party tools.

- Time-to-Value and Agility:** How urgently is results needed? Buying a solution typically yields immediate ROI: set up a subscription and start using it within weeks. Building requires up-front development. Analysts advise *“buy first, deliver something visible, fast, and safe. Then decide if it’s worth building”* ([2] www.zartis.com). Early wins build stakeholder support and avoid stagnation. However, if a company has very long planning cycles (as in some large pharmas), a longer build timeline may be tolerated if properly funded.
- Cost and Budget Constraints:** Total cost includes not just upfront but multi-year TCO. Project budget size is a factor: small projects (\$100K–500K) might lean towards buying, whereas multi-million-dollar platforms might be developed. In general, commercial projects with well-established vendors often cost an order of magnitude less to start than in-house builds ([4] intuitionlabs.ai). That said, one must compare license fees over time vs amortized development cost: for very long-lived projects, building might break even or become cheaper if the vendor would otherwise charge heavy perpetual fees.
- Data Availability and Quality:** Building a custom model requires clean, well-curated data. Pharma data is notoriously heterogeneous and siloed ([25] www.pharmexec.com). If data resides across disparate systems with no interoperability, building an AI from scratch is harder. Conversely, vendors may offer data integration expertise or even pre-trained models on broader datasets. For example, a drug-discovery vendor might have aggregated cheminformatics databases that augment your internal data. Trustworthiness and lineage of data (regulatory provenance) are also vital.
- Regulatory and Privacy Compliance:** AI handling patient or trial data must comply with HIPAA, GDPR, FDA/EMA guidance, etc. Buying from a vendor means ensuring they also comply fully; any breach or audit issue may still implicate you. Building in-house means total control, but also all the compliance burden. In either case, thorough due diligence is needed. Some vendors specialize in life sciences and already integrate audit trails and governed workflows.
- Vendor Landscape and Maturity:** Does a specialized solution already exist? If the function is common (e.g. medical image analysis, supply forecasting, pharmacovigilance triage), there may be mature products. Commercial platforms (e.g. for clinical trial optimization like Medidata’s Predict™, or digital marketing like Veeva CRM) accelerate adoption. However, be wary of “overachieving” vendors that oversell immature tech. Vendor viability and domain expertise should be vetted via pilots and references.
- Talent and Capabilities:** Do you have or can you hire AI experts and engineers? Building demands a strong multidisciplinary team: data scientists, ML engineers, domain scientists, DevOps, plus project management. If this talent is unavailable or expensive to retain, buying a solution leverages the vendor’s team. Even after purchase, you still need people to operate and interpret the system, but not build it from the ground up.
- Scalability and Maintenance:** AI systems need constant upkeep (data drift monitoring, model retraining, scaling to growing users). Vendors usually handle much of this transparently. In-house builds need an internal MLOps roadmap. Hidden costs like security patches, cloud infrastructure scaling, and retraining pipelines can dominate expenses beyond the first year ([24] intuitionlabs.ai).
- Intellectual Property (IP) and Future Value:** In-house AI can become IP that adds to company valuation and competitive moat. A homegrown algorithm that wins significant clinical trials or sales could be influential. With vendor tools, all IP typically resides with the vendor (except any specific data). In fast-moving fields, owning the know-how can pay off, but only if the systems are indeed distinctive. If the advantage would likely be short-lived (as AI becomes ubiquitous), then the IP benefit may not justify the cost.

These considerations often pull in opposite directions. The smart approach is often **hybrid**: begin with buying to learn and prove value, then shift to building once the use-case is validated and the business case is clear. As one analysis recommends, “start by buying to achieve quick wins and validate use cases, then build custom systems where proprietary competitive advantage truly requires it” ([1] intuitionlabs.ai). Table 2 below outlines some typical pharma AI applications and suggests whether they are usually built or bought.

Domain	AI Applications	Examples / Solutions	Typical Approach
Drug Discovery	Molecule screening, target ID, protein/structure prediction	DeepMind’s AlphaFold2 (protein folding model) ([10] moneyweek.com); Insitro’s ML on biological markers ([11] apnews.com)	Vendor-built models; hybrid integration
Drug Development	Clinical trial design optimization, patient matching, RWE analysis	Formation Bio (AI for trial admin) ([12] time.com); AI patient recruitment algorithms	Often vendor/partnership for trials
Manufacturing & QC	Predictive maintenance, process optimization	Siemens’ AI in pharma manufacturing; in-house data analytics	Both (often build analytics platform)
Sales & Marketing	CRM automation, HCP targeting, content personalization	Workbooks CRM with AI tools ([14] www.techradar.com); Veeva CRM Predict; Salesforce Einstein	Primarily buy (CRM vendors, SaaS)
Patient Engagement	Adherence coaching, chatbots, digital therapeutics	UpDoc insulin-management assistant ([26] www.axios.com); patient chatbot platforms	Typically vendor or joint startup

Domain	AI Applications	Examples / Solutions	Typical Approach
Supply Chain & Ops	Demand forecasting, inventory planning, logistics	SAP IBP with AI; third-party supply-chain AI analytics	Often mix (industry SaaS platforms)

Table 2. Examples of AI use-cases in pharmaceutical R&D and commercial domains, with typical implementation approaches. (Sources: Table rows are illustrative; see cited examples.)

The table illustrates that **R&D-focused tasks** (discovery, modeling) often leverage specialized AI platforms or startups with deep domain expertise. For example, AlphaFold2 is a public model whose predictions are now used by many, rather than each firm replicating protein-folding code (^[10] moneyweek.com). Similarly, drug screening may use vendor tools (or in future, large generative-model APIs trained on chemical data). By contrast, **sales and marketing applications** tend to use standard CRM and analytics solutions, so are usually purchased. Patient-facing tools are often developed in partnership: for instance, UpDoc’s AI assistant was a startup funded by Mayo Clinic and Eli Lilly (^[26] www.axios.com). Supply-chain optimizations may be delivered by footsteps of enterprise software providers.

Ultimately, **the decision framework** should consider the specific context of each initiative: Is this something we *absolutely need to own*, or is it better to plug into an existing solution? Sections below provide deeper analysis in various aspects, with data and expert perspectives to back them up.

Costs, Total Cost of Ownership, and Economic Analysis

A critical dimension of build-vs-buy is economics. Beyond one-time development or procurement costs, companies must consider ongoing expenditures and ROI. Studies suggest that in-house AI projects in life sciences demand a *substantial* financial commitment. For example, one analysis estimates a custom enterprise AI project might require roughly **\$500K–\$2M** to launch, plus **30–40%** of that amount each subsequent year for maintenance (^[4] intuitionlabs.ai). Conversely, subscribing to commercial AI platforms can cost on the order of **\$50K–\$200K** up front, with smaller annual fees (10–20%) (^[4] intuitionlabs.ai). In other words, building can be an order of magnitude more expensive initially. (These figures align with figures for related AI staffing and infrastructure: a single senior AI researcher can cost \$300K–\$500K per year (^[27] xenoss.io).)

Pharma finance teams should also consider the **opportunity cost**. Tying up, say, \$2M of capital and a year of effort on building a model means those resources are not spent elsewhere. If a ready-made solution can deliver 80% of the value for 20% of the cost, the remaining 20% of custom value must be strategic to justify building. Moreover, hidden costs—data cleaning, security, regulatory auditing—tend to accumulate in internal projects over time; reports warn that after 9–12 months, the majority of AI spending can go to these unplanned elements (^[24] intuitionlabs.ai).

An integrated financial model should include:

- **Development costs:** salaries, cloud/compute, data engineering, UI/UX, validation.
- **Licensing costs:** vendor fees, seats, modules.
- **Maintenance:** bug fixes, retraining, downtime risks.
- **Opportunity cost:** what other R&D or commercial projects could be funded.
- **Value creation:** expected revenue uplift or cost savings from AI (often speculative).

Some vendors or consultants provide simple calculators. For instance, Xenoss (an AI advisory firm) published a comparison (Table 1) contrasting build vs buy, reproduced in Table 1 above (^[16] xenoss.io). Also, a helpful approach is to compute a **breakeven horizon**: at what point cumulative costs of buying exceed building? This depends on usage scale. If a model will be used by many products or markets for a decade, building software might amortize; but if it’s a one-off analysis, buying is usually cheaper.

Integration, Data, and Technology Architecture

Building AI in pharma often means integrating with complex legacy systems. Pharmaceutical R&D and commercial data are typically **silod**: separate databases for lab experiments, clinical trials, electronic health records (EHR), supply logistics, marketing, etc. Legacy incompatibilities (different identifiers, data formats, standards) make data integration challenging (^[25] www.pharmexec.com). A generic vendor model may not easily plug into these silos without heavy ETL (extract-transform-load) work. Internal builds must also tackle this ETL themselves.

- **Cloud vs On-Premises:** Another build-or-buy variant is *where* to run AI. Many companies hesitate to put sensitive data in public clouds. If on-prem, building means procuring or expanding data center capacity for GPUs, which can cost millions per year (^[17] intuitionlabs.ai). Buying often means cloud SaaS; while it trades capital expense for operating expense, pharma must ensure cloud compliance (e.g. FedRAMP, HIPAA).
- **Model Lifecycle:** Commercial solutions may abstract the model lifecycle (e.g. automated retraining), but custom AI requires building an MLOps pipeline. As Sheikh Sharjeel (AI strategist) notes, non-obvious costs like CI/CD pipelines, orchestration, security, and governance can dominate the TCO after the first year (^[24] intuitionlabs.ai). Companies must plan for continuous monitoring (model drift, bias checks) if they build a system that will be used long-term.
- **Adaptability:** If business requirements will change rapidly (e.g. new diseases, new regulatory mandates, new marketing channels), a build must be adaptable. Vendors frequently update features and models to keep up with trends (e.g. adding GPT-style capabilities). However, vendor roadmaps are out of your control; buying locks you into their pace of innovation.
- **Hybrid Approaches:** Teams sometimes adopt partial build: e.g. using a vendor's core AI engine but developing proprietary data-processing modules around it. Or building a custom UI/workflow on top of a purchased API. These hybrid architectures attempt to balance speed with control.

Given the above, integration complexity often favors buying for **horizontal needs** (common data prep, analytics) and bundling multiple functions. In contrast, if the use-case is narrow and data well-structured, a build might be more viable. For example, a pharmaceutical manufacturer might develop an in-house AI for a specific process analytics tool, whereas supply-chain forecasting might leverage an industry SaaS.

Organizational and Talent Considerations

AI initiatives fail or succeed based as much on people and process as on algorithms. The **organizational readiness** is crucial:

- **AI Expertise:** Building AI requires data scientists, ML engineers, DevOps specialists, and domain scientists all collaborating. According to industry sources, specialized AI talent commands very high compensation (e.g. \$200K–\$500K for data scientists or principal researchers (^[27] xenoss.io)). Very few pharma companies have enough such talent internally. In buying, the burden shifts: you need only integration engineers and business-side users.
- **Interdisciplinary Teams:** For a custom build, teams must also include regulatory and domain experts to ensure compliance. A cross-functional steering committee often helps align goals. In some partnerships, hiring an internal AI product manager to coordinate vendor and biz is key.
- **Internal Alignment:** Implementing an AI system (built or bought) requires organizational change management. A custom build may require retraining of staff on new workflows and systems. Vendor tools often come with user support and pre-built interfaces, which can ease adoption (^[28] intuitionlabs.ai). Pharma users (e.g. biologists, clinicians) may be skeptical of black-box models, so governance and explainability can be more easily addressed if the solution is homegrown—or conversely, vendors may provide explanation modules. Either way, user trust is earned through transparency, and that factors into build vs buy: who can provide sufficient trust signals?
- **Corporate Culture and Strategy:** As Takeda's CEO Christophe Weber remarked, his company decided **six years ago** to become "data-driven, tech-driven"; they now use AI across R&D and manufacturing (^[29] www.axios.com). Culture shifts like this take time. If leadership is cautious about AI, they may prefer proven external tools. Conversely, frontrunners may greenlight ambitious in-house labs.

In sum, companies must ask: *Do we have or can we acquire the skills to build this well?* If not, partnering with an AI vendor or consultancy often dramatically raises success odds. Indeed, many life sciences firms form partnerships or joint ventures (a hybrid route) to blend vendor expertise with internal data access. The approach of “**build and buy both in the right order**” may involve hiring a small AI team to scope requirements, deploying a vendor proof-of-concept, then gradually insourcing components that show the greatest value (^[2] www.zartis.com) (^[1] intuitionlabs.ai).

Case Studies and Industry Examples

Eli Lilly and Nvidia (Build / Partnership): In 2025, NVIDIA's CEO Jensen Huang highlighted Lilly as a pharma pioneer in AI. Lilly announced a multi-year partnership with NVIDIA to build a dedicated supercomputing infrastructure for “scientific AI agents” that will design experiments and accelerate research (^[30] www.axios.com). This is a textbook *build* approach at scale: Lilly is investing in bespoke hardware and software to own a new AI-driven R&D engine. The strategic bet is that such an in-house “digital lab” will yield faster discovery (potentially helping to reach NVIDIA's \$1T prediction and expanding Lilly's pipeline).

Bayer Pharmaceuticals (Partnership / Internal Build): Bayer (U.S.) COO Sebastian Guth described AI as a tool for “driving workflow and operational efficiencies” and that AI at scale could help develop medicines “that would have otherwise likely not seen the light of day.” (^[31] www.axios.com) Bayer invests in AI to augment human ingenuity in drug discovery. While Bayer has its own data science teams, it also partners with technology firms (for example, Bayer has previously collaborated with DreamWorks and BenevolentAI on drug screening). Bayer's strategy appears hybrid: using vendor and academic tools to accelerate R&D, while internal teams integrate them into the pipeline.

Takeda (Internal Transformation): Takeda CEO Christophe Weber emphasized that Takeda decided years ago to make AI a core pillar, and now applies AI “to accelerate research and clinical development,” including predicting drug effects and optimizing trial enrollment (^[29] www.axios.com). Takeda, a large global pharma, invests in building data infrastructure internally. For instance, Takeda has recruited hundreds of data scientists and formed an internal “AI Center of Excellence.” They also acquired AI startups (e.g. gene therapy company Maverick partnered, plus collaborations with IBM Watson Health in the past) to combine buy/build. The lesson is that a company with enough scale can justify building broad AI platforms, though even Takeda often taps external partners for specialized tasks.

Insitro and Pharma (Outsourced AI Services): Insitro demonstrates a biotech-as-vendor model: it built proprietary AI/ML technology and offers it as a service to pharma. Having itself invested in ML talent and computing, Insitro's customers (Lilly, BMS, etc.) essentially *buy* decision support. The output (e.g. target predictions) become inputs into the pharma's internal pipeline. This leverages Insitro's deep learning expertise without requiring each partner to build the same models.

Customer Engagement (Buy): On the commercial side, solutions are often bought. For example, many sales teams now use AI-enhanced CRMs rather than in-house systems. The Workbooks CRM (with newly integrated AI modules) was rapidly adopted by early pharma customers, because it plugs into existing sales workflows (^[14] www.techradar.com). Few companies are building their own CRM platforms because vendor products (Salesforce/Veeva, IQVIA, etc.) are robust and validated. Similarly, AI chatbots for patient support (like UpDoc's insulin-DM bot) were not developed by pharma firms themselves but by startups funded by industry players (^[26] www.axios.com).

Data Quality (Challenge for In-House AI): Separated from build/buy, the quality of underlying data is a major determinant. As one report notes, a Phase III clinical trial now generates millions of data points (≈6 million on average, up from 1 million a decade ago) – much of it from external sources like electronic health records and wearables (^[32] www.pharmexec.com). However, these data are often incomplete or biased, and integrating them is like stitching together a fragile mosaic (^[25] www.pharmexec.com). A purchased AI platform may come with pre-built connectors to common clinical databases or standards (e.g. FHIR), whereas an in-house solution requires building those integrations from scratch. Thus, some pharma companies first invest in data warehouses or lakehouses (as a build project) before even tackling AI on top of that.

Regulatory, Ethical, and Risk Considerations

Pharma is a **highly regulated** industry. Any AI used in R&D or commercial activities must contend with data privacy, patient safety, and compliance frameworks:

- **Privacy and Data Governance:** Clinical and patient data fall under GDPR, HIPAA, and other laws. Off-the-shelf AI tools (especially from big tech) may not be compliant by default. Buying an AI tool for analysis of patient data requires ensuring the vendor's handling protocols are approved. Building internally gives you full control over encryption, anonymization, and audit trails, but also full responsibility. In either case, data governance plans (audit logs, data retention policies) must be part of the project scope.
- **Regulatory Approval of AI-Driven Products:** If the AI itself is used to make therapeutic decisions (e.g. an AI that selects trial patients or recommends a treatment), it may be subject to FDA/EMA regulation. Agencies are beginning to issue guidance on "Software as a Medical Device" and AI. Regulators generally require documented validation: showing the AI works as intended. A major potential risk of buying a "black-box" AI is that it may be difficult to fully validate its every mode of operation. Conversely, building from scratch means investing in compliance documentation (validation plans, risk analyses, change controls) from day one.
- **Intellectual Property and Data Sharing:** Using external AI tools often means license agreements over intellectual property. Some vendors may claim rights to improvements or derivative works. Pharma must be cautious not to inadvertently license out insights. Building in-house avoids such IP entanglement but requires internal IP management (e.g. patenting algorithms).
- **Algorithmic Bias and Explainability:** Within pharma, bias can be especially dangerous (e.g. underrepresenting minority groups in data). Commercial AI systems may be trained on datasets that are not fully open to auditors. In contrast, an in-house AI can be explicitly tested and biases addressed. Both paths need governance: companies must decide how to document AI behavior. Some firms mitigate this by choosing AI platforms that support explainability modules (SHAP values, confidence scores) or by developing in-house explainable models even if core models are vendor-provided.
- **Business Continuity and Vendor Risk:** Relying on a vendor carries vendor-risk: stability of the provider, business continuity, and support responsiveness matter. If the vendor goes bankrupt or changes strategy, your AI tool could vanish or degrade. In-house systems avoid that dependency, at the cost of having sole responsibility for uptime.

Given these factors, regulatory and risk concerns often tip the balance toward buying in non-differentiating areas, because specialized AI-for-health vendors typically have robust compliance features built in. But for mission-critical or safety-critical systems, the extra assurance of controlling the stack (via building) can be worth the overhead.

Discussion of Implications and Future Directions

Looking ahead, several trends will influence build-vs-buy decisions in pharma AI:

- **Maturation of AI Ecosystem:** As generative AI and foundation models mature, the line between build and buy may blur. Companies might increasingly *subscribe to large language or vision models (via APIs) and tune them themselves*. For example, a pharma could fine-tune an open-source LLM on its own medical literature: this is partly buying (base model from third party) and partly building (fine-tuning and specialized prompt engineering). This approach lowers the in-house training barrier. However, heavy reliance on external model providers (OpenAI, Google, etc.) raises new vendor-risk and data privacy issues.
- **Data Sharing and Industry Platforms:** Some future scenarios involve industry-wide platforms. Partnerships like the Pistoia Alliance or academic consortia may provide common AI resources. Pharma consortia (with anonymized data) could build shared AI tools. If successful, this collective build effort could reduce individual need to build. Conversely, proprietary data silos will keep some companies wishing to build private systems insulated from such shared platforms.
- **Generative AI and Novel Modalities:** Emerging AI applications (e.g. generative peptide design, automated lab robotics) could become high-value. Firms might try to build novel generative models that propose molecules – though today most use external research. If an AI-designed therapy emerges, owning that generative engine's IP would be invaluable, so building (or partnering in co-development) could be prioritized in R&D. Similarly, "AI Lab Assistants" (autonomous lab robots) might be deployed; current examples like Reverie Robotics or Benchling science operations hint at this. Where such an evolution occurs, early adopters with in-house expertise gain an edge.

- **Regulatory Guidance for AI:** The regulatory landscape is evolving. New FDA/EMA guidelines on AI/ML in drug development will create clearer expectations. If regulators start expecting robust validation of AI tools (and possibly approval of certain algorithms), pharma will need either very stable vendor relationships or the capacity to build compliant solutions. In some cases, regulators might encourage use of qualified AI tools (like requiring use of certain validated models as industry standard).
- **Cost Reduction and Commoditization:** Over time, as AI becomes more common, prices for commercial solutions may drop or become subscription-based by default, making buying even more attractive for standard use-cases. A few years from now, querying a powerful generative model for chemistry could cost only a few dollars; building a similar capability would be relatively wasteful except for unique customization. The emphasis will thus be on *differentiation*. Companies will ask: “Will this in-house AI truly beat what any other company could buy off the shelf?”
- **Strategic Shifts in Pharma Structure:** Bigger pharma may spin off separate digital/AI arms (as Glaxo did with their digital innovation team) to experiment rapidly (a de facto build approach), while smaller biotechs might rely entirely on partnerships and contract research. Mergers and acquisitions may increasingly include AI talent as the asset (e.g. acquiring an AI startup along with its products).

The consensus from industry experts, as captured in recent surveys and commentaries, is that **successful AI adoption is less about the technology and more about execution** (^[9] www.pharmexec.com). Merely buying a product or starting an ambitious build without organizational alignment often leads to failure. The difference between “AI aspirational” and “AI advanced” companies, a PharmExec study notes, lies in broad strategy, governance, and cross-functional integration – not just investment level (^[9] www.pharmexec.com). In practice, companies that treat AI as a core strategic initiative (with dedicated budgets, cross-department committees, and executive sponsorship) tend to navigate build-vs-buy in a more structured way.

Conclusion

Building versus buying AI in the pharmaceutical sector is not a binary choice but a strategic continuum. Each R&D project and commercial program demands a tailored approach. The evidence suggests a **pragmatic hybrid path**: begin with vendor platforms to test concepts rapidly and with minimal investment, then evaluate where winning requires bespoke innovation. As Dr. Laurent and colleagues recommend, “*plan a sequence — know when to buy, when to build, and how to orchestrate both to scale sustainably.*” (^[33] www.zartis.com).

Key takeaways for pharma leaders include: identify **core applications** where owning AI is critical; assess total cost and timeline carefully; understand that generically available solutions can often cover 70–80% of needs; and remember that vendor solutions can themselves become standards (so long-term competitive advantage may reside elsewhere). Align the build/buy decision with overall AI maturity: if the organization is just starting out, prioritize fast value via purchase; if it's mature and sees a gap in commercial offerings, consider investing in build. And always evaluate the **risk-adjusted ROI**: if a custom AI provides materially higher drug success probability or marketing lift beyond what an average vendor can, it may justify the extra cost (^[4] intuitionlabs.ai).

In the years ahead, no single path will apply universally. Some pharma giants will erect their own AI-centered “labs” and toolchains, while others will plug into a growing marketplace of AI-as-a-service offerings. The competition in pharma is high-stakes and time-pressured: mistakes in R&D cost companies even more than in other sectors. Therefore, the build vs buy framework outlined herein – grounded in cost data, industry cases, and strategic analysis – should be revisited constantly as technology and competitive landscapes evolve. By combining rigorous evaluation with flexible strategy, pharma R&D and commercial teams can harness AI effectively, accelerating innovation and ultimately delivering better medicines to patients.

External Sources

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.

© 2025 IntuitionLabs.ai. All rights reserved.