Generative Antibody Design: Chai-3 and Pharma AI Strategy

6/13/2026 • 30 min read

generative antibody design | chai-3 model | ai drug discovery | zero-shot design | pfizer ai strategy | pharma ai playbook | de novo protein design | biotechnology

Executive Summary

In June 2026, Pfizer announced a landmark licensing agreement with AI startup **Chai Discovery**, granting the pharmaceutical giant access to Chai's advanced generative antibody design platform, including the novel **Chai-3** model, as well as a bespoke AI model trained on Pfizer's own proprietary data ([1] www.biospace.com) ([2] www.businesswire.com). Chai-3, previously undisclosed, represents a "step-change" in AI-driven antibody generation, **doubling the success rate** of its predecessor (Chai-2) and producing candidate antibodies that meet therapeutic standards, including multispecific formats and binders for traditionally "hard" targets ([2] www.businesswire.com) ([3] qz.com). The so-called "zero-shot" design capability of this platform means it can generate functional antibody sequences directly from the target antigen and epitope, **without requiring target-specific training data**. In laboratory tests, Chai-2 achieved roughly a 16–20% hit rate (validated binding) across dozens of previously unaddressed targets ([4] finance.yahoo.com) (colab.ws) – over **100-fold better** than prior computational methods (<0.1%) ([4] finance.yahoo.com). Pfizer's deal – which includes early access to Chai-3 and a custom AI pipeline built on Pfizer's own datasets – exemplifies the broader industry shift: major drugmakers are increasingly moving from AI *proof-of-concept* research into concrete **build-vs-partner** strategies for real-world drug discovery workflows ([2] www.businesswire.com) ([5] www.techtarget.com). Indeed, recent months have seen massive industry partnerships (e.g. Merck's $1B deal with Google Cloud, Novo Nordisk's alliance with OpenAI, and Eli Lilly's collaborations with NVIDIA and Insilico) aimed at accelerating R&D with generative AI ([6] www.fiercebiotech.com) ([5] www.techtarget.com). This report analyzes the technical breakthroughs behind Chai's platform—especially its zero-shot generative design—examines the strategic implications of Pfizer's license (including the custom model), and situates this case in the **build-vs-partner playbook** that pharma companies face when adopting frontier AI. We draw on extensive data, expert commentary, and case studies to evaluate the current landscape, the rationale for partnership vs in-house approaches, and the future directions of AI-powered drug discovery.

Introduction and Background

Drug discovery has long been characterized by **high costs, long timelines**, and low probabilities of success. Empirically, the "space" of possible antibody sequences (estimated on the order of 10^{18} variants) vastly exceeds any feasible experimental search ([7] pmc.ncbi.nlm.nih.gov). Traditional antibody development (animal immunization, phage/yeast display) explores only minute fractions (~10^{11} variants in large libraries) ([7] pmc.ncbi.nlm.nih.gov). Consequently, discovering high-affinity, developable antibodies against novel targets has been time-consuming and hit-driven. Even sequences suggested by computational methods historically showed **hit rates below 0.1%** in experimental validation ([4] finance.yahoo.com).

In recent years, **deep learning and generative AI** have promised to transform this landscape. AI-driven protein design aims to "learn" from vast data on antibody structures and antigen interactions to predict new functional sequences. Notable advances include *structure-prediction* models (e.g. DeepMind's AlphaFold, RoseTTAFold) and *language/sequence models* (e.g. prot-level LLMs). However, until lately these methods mainly optimized known scaffolds or affinities, still requiring deep wet-lab screening. The breakthrough Chai Discovery announced in mid-2025 – its *Chai-2* model – marked the first demonstration of **fully de novo ("cold start") antibody design with high efficiency** ([4] finance.yahoo.com) (colab.ws). Chai-2, a multimodal generative model, accepts as input only the target antigen and epitope (structure) and generates entirely novel complementarity-determining regions (CDRs) for antibodies or nanobodies. In a broad test against 52 diverse targets (none of which had a known binder), Chai-2 produced at least one validated hit for ~50% of targets, often achieving strong (nanomolar) affinities and good developability (colab.ws) ([8] finance.yahoo.com). This zero-shot approach (no target-specific tuning) achieved **16% hit rate** in under two weeks of design and testing (colab.ws) ([9] finance.yahoo.com). By contrast, earlier computational platforms typically required screening thousands of candidates per target and yielded near-zero success ([4] finance.yahoo.com). As Chai co-founder

Joshua Meier put it, Chai-2 is “like Photoshop for proteins” – enabling precise, rapid, and intuitive design of biologic therapeutics ([10] finance.yahoo.com).

The importance of such generative capability is well documented. Meng et al. note that machine learning can capture the high-dimensional distribution of known antibodies and infer new sequences rapidly ([11] pmc.ncbi.nlm.nih.gov). In effect, deep generative models “detect high-order amino acid interactions” and can explore sequence space far faster than random or rational designs ([11] pmc.ncbi.nlm.nih.gov). (For example, Chai-2’s designs included CDR loops “distinct from [natural] VHHs… indicating substantial generalization beyond the training dataset” ([12] pmc.ncbi.nlm.nih.gov).) The net result is that AI methods can navigate the 10^{18}-size search space without exhaustive screening ([7] pmc.ncbi.nlm.nih.gov).In practical terms, Chai-2 allowed researchers to go “from in silico generation to lab validation in under two weeks” ([9] finance.yahoo.com), versus months or years needed by traditional methods.

At the same time, the pharmaceutical industry faces an intensifying imperative to trim R&D costs, confront patent expirations, and fight demographic and competitive pressures. Leading companies report multi-billion-dollar revenue quarters while also reinvesting heavily into their pipelines. For instance, Pfizer reported Q1 2026 revenues of $14.45 billion (driven by oncology growth and new launches) and actively advanced multiple clinical assets ([13] qz.com). To maintain this momentum, firms are leveraging every tool – and AI is viewed as a critical accelerator. Already by early 2026, major pharmas are shifting from experimental pilot projects to enterprise-scale AI deployments. Examples abound: AI agents are automating regulatory writing, manufacturing QA, and clinical logistics; meanwhile, generative models are tackling drug design. This report focuses on the latter: the pairing of Pfizer with Chai’s antibody-design AI is emblematic of how generative “foundational” models (akin to upgraded versions of LLMs for biology) are being integrated into pharma R&D pipelines.

In the sections that follow, we delve into the technological breakthroughs of Chai’s platform (Chai-2 and Chai-3), examine the specifics and strategic rationale of the Pfizer partnership, and analyze the broader strategic question for pharma: **should a company build its own AI solutions or partner with specialized AI vendors (or do both)?** We draw on a mix of technical literature, press releases, expert commentary, and reported statistics to present a comprehensive picture of the current state and future trajectory of AI in drug discovery.

Generative AI for Antibody Design

The Biology and the Innovation

Antibodies are large, immune-system proteins that bind specific antigens with high affinity. Therapeutically, monoclonal antibodies have become one of the most successful drug classes (177 approved antibodies on the market as of mid-2024 ([14] pmc.ncbi.nlm.nih.gov)). However, the fundamental challenge remains: finding an antibody sequence that binds a chosen target and has drug-like properties. The space of possible antibody sequences is astronomical (~10^{18} variants of CDR loops ([7] pmc.ncbi.nlm.nih.gov)), so empirical methods (immunizing animals or screening phage display libraries) only sample a tiny fraction. Historically, such methods required *millions* of dollars and months to years to identify a single lead. Even rational computational approaches (Rosetta-based design, motif grafting, or earlier machine learning models) have struggled: hit rates were typically well below 1%, and often zero without extensive follow-up screening ([4] finance.yahoo.com).

Generative AI aims to **learn the underlying patterns of antibody-antigen interactions directly from data**. Modern approaches include sequence-based LMs (trained on millions of protein sequences), structure-based optimization, and hybrid methods. For instance, the recent literature has explored conditional language models that infill antibody sequences (e.g. IgLM ([15] pmc.ncbi.nlm.nih.gov)), latent diffusion models conditioned on antigens, and designs guided by energy-based scores. Google/DeepMind and others have advanced protein structure prediction, enabling better evaluation of candidate designs. But the ultimate goal – a “zero-shot” antibody generator – implies a model that can

create truly novel, high-quality binders for an arbitrary target epitope without ever having seen similar binders. Chai-2 (*Zero-shot antibody design in a 24-well plate*, Boitreaud et al., bioRxiv 2025) is one of the first such systems. By feeding Chai-2 a 3D structure or epitope patch of an antigen, it autogenerates complete CDR sequences that form a functional antibody. In lab tests across 52 diverse antigens, **Chai-2's hit rate was ~16%**, i.e. almost double-digit hits out of only ~20 designs each on average (colab.ws) ([4] finance.yahoo.com). Specifically, about half the targets yielded at least one validated binder (colab.ws). (For comparison, a halving of failure rate effectively means success rate doubled relative to earlier benchmarks ([3] qz.com).) Chai-2-produced binders also showed desirable properties – nanomolar affinities, specificity, and good developability – demonstrating that the generative model learned not only to bind but to satisfy drug-like constraints ([8] finance.yahoo.com). The same system achieved ~68% wet-lab success in designing minibinders (small non-antibody proteins) with picomolar binding (colab.ws), highlighting its broad potential in molecular design.

Importantly, Chai-2 achieved these results with "under two weeks" from computational output to validated lab hit ([9] finance.yahoo.com) (colab.ws). By contrast, traditional antibody development often requires *months of iterative screening and maturation*. This compression of the discovery loop is transformative. As one biotech leader toasted, in one anecdote Chai-2 solved a previously intractable antibody problem (that had consumed $5M R&D) in mere hours, with lab validation in under 2 weeks ([16] finance.yahoo.com). In the words of Mikael Dolsten (Pfizer's ex-CSO), Chai-2 "holds high potential for de novo design of medicines with short turnaround times" ([17] finance.yahoo.com).

From Chai-2 to Chai-3: Next Generation Performance

Building on Chai-2's success, Chai Discovery has since developed **Chai-3**, a more powerful generative model. Unlike Chai-2 (which was publicized in mid-2025), Chai-3 remained under wraps until the Pfizer deal. According to Chai's announcement, Chai-3 delivers a **"step-change" improvement in antibody design** ([2] www.businesswire.com). It "doubles the success rate" of Chai-2 and yields candidates meeting therapeutic quality thresholds ([2] www.businesswire.com). Operationally, this implies that if Chai-2 hit ~16–20%, Chai-3 hits may approach ~35–40% on similar problems (though exact figures are not yet disclosed). Further, Chai-3 is said to enhance capabilities in designing multi-specific antibodies, hitting "hard-to-drug" targets, and generalizing better across diverse antigens ([2] www.businesswire.com). In plain language, Chai-3 can solve a wider array of challenging binding problems with higher reliability.

The license announcement emphasized that Pfizer will have among the first access to this model ([2] www.businesswire.com) ([1] www.biospace.com). What is known of Chai-3's architecture suggests it likely uses more advanced training (possibly larger datasets and parameters) and improved modeling of binding energetics. For example, Chai-2 already integrated both sequence and structural information; Chai-3 may refine this with deeper networks or hybrid diffusion frameworks. The result is akin to going from "AlphaFold" to "AlphaFold2" in antibody space. The press release highlights that Chai-3 not only improves binding affinity but also expands design space (e.g. engineered multi-domain formats) ([2] www.businesswire.com).

Crucially, Chai-3 retains the **zero-shot design** principle of Chai-2: it generates full CDR sequences de novo from the target without target-specific retraining ([2] www.businesswire.com) ([18] www.biospace.com). This is a generative leap: some computational design methods (like RFdiffusion) require starting from an existing scaffold or framework (and even then devote weeks to design and yeast display) ([12] pmc.ncbi.nlm.nih.gov). Chai-3, by contrast, was built to generate entirely new antibodies much faster. At its core, it functions as a "generator network" guided by target structure. While details are proprietary, one can infer that Chai-3's training likely included orders-of-magnitude more antigen-paired data or innovative loss functions to drive therapeutic quality.

Evidence of Chai-3's impact comes from Pfizer's own statements and media analysis. The Quartz article notes that Chai-3 halves the failure rate of antibody design compared to Chai-2, while still producing "candidates that meet therapeutic standards" ([3] qz.com). In other words, Chai-3 achieves higher positive rates without sacrificing developability or binding strength. We should also note that the Pfizer license includes not just run-of-the-mill Chai-3, but a **"custom**

model" tailored to Pfizer's data and workflows ([1] www.biospace.com) (discussed below). This implies Pfizer can selectively train or fine-tune generative outputs using its proprietary antigen/antibody datasets, further boosting performance for their programs.

Finally, it is instructive to place Chai-3 in context with academic advances. For instance, Baker's 2025 *Nature* study demonstrated that a diffusion model (RFdiffusion) could design antibodies with atomic accuracy, validated by cryo-EM and binding assays ([12] pmc.ncbi.nlm.nih.gov). Those designs took ~6 weeks and extensive lab work. Chai-3's claim to rush this cycle to under a fortnight with doubled yield is striking. In summary, Chai-3 represents one of the most advanced generative platforms for biologics to date, promising to reduce the time and effort to find leads by **orders of magnitude**.

Pfizer's License of Chai-3 and the Custom Model

Details of the Agreement

On June 4, 2026, Chai Discovery announced that it has signed a **license agreement** with Pfizer to integrate Chai's AI platform into Pfizer's drug discovery engine ([1] www.biospace.com) ([2] www.businesswire.com). Under this deal, Pfizer will gain *early access* to Chai-3 – the latest model – as well as tooling and software for generative design. Critically, the license includes provision for a **custom AI model trained on Pfizer's own data**, tailored to Pfizer's internal processes ([1] www.biospace.com). In practice, this means Pfizer can combine Chai's cutting-edge generative engine with the wealth of proprietary data it has (including structures, screening results, sequence-function pairs, etc.) to fine-tune or adapt outputs. Joshua Meier (Chai co-founder) emphasized that this collaboration puts "Chai's software directly into the hands of one of the world's leading drug discovery organizations," leveraging Chai's AI with Pfizer's scientific depth ([19] www.biospace.com).

The press materials highlight two key outcomes: (1) Pfizer will be among the first to deploy Chai-3 for real target programs ([2] www.businesswire.com); (2) Pfizer will receive a model version that incorporates its proprietary knowledge. This custom model likely means Pfizer can either provide their own data to retrain/finetune a Chai model, or possibly co-develop the model architecture with Chai, under license terms. Although financial terms were not disclosed, such deals typically involve licensing fees and possibly milestone payments tied to research outcomes (akin to biotechs licensing discovery tools). For context, many big pharma have announced multi-million or multi-$B AI deals in 2025–26 (see below), so it is reasonable that Pfizer's deal is significant.

The **strategic rationale for Pfizer** is clear. Antibody therapeutics remain a core of Pfizer's biologics portfolio, and new modalities (bispecifics, cell therapies) rely on antibody engineering. By partnering rather than building the capability from scratch, Pfizer can leap ahead in generative design. As Chad—is known, building an AI platform internally demands assembling vast datasets, hiring AI specialists, and enduring long development times ([20] intuitionlabs.ai) ([21] intuitionlabs.ai). The Chai partnership sidesteps much of that initial investment. Pfizer still gains custom integration (through the bespoke model), but offloads model R&D to Chai's team of AI veterans from OpenAI, Meta, etc. As Meier put it, this combines "Pfizer's scientific depth, data and discovery capabilities" with Chai's platform ([19] www.biospace.com). In essence, Pfizer has chosen to **partner** on AI: leveraging external innovation on generative models while contributing internal assets (data, targets) to co-create value.

This arrangement contrasts with pharma firms that attempt purely internal development. For example, some companies have built AI centers or alliances (Novartis's AI lab with Microsoft in 2019; Roche/Genentech doing internal ML efforts), but even these often eventually partner for specific tools. In this case, Pfizer opts for a partner/licensing route, which was prescient enough to secure "a previously undisclosed model". It demonstrates confidence in Chai's technology and a willingness to integrate a startup's software into Pfizer's workflows. The agreement thus marks a significant milestone: it

effectively **places generative AI at the heart of Pfizer's biologics R&D engine**, ready for use on current or future antibody projects.

Custom Model and Proprietary Data

A unique aspect of the deal is the **custom model** that leverages Pfizer's proprietary data ([1] www.biospace.com). This likely reflects a hybrid strategy: Pfizer will not simply run Chai-3 "out-of-the-box", but will refine it with internal resources. In practice, Pfizer could provide large amounts of antigen–antibody binding data, structural complexes, and phenotypic outcomes to train Chai's foundations. Combining a general antibody design model with Pfizer-specific finetuning could dramatically improve performance on Pfizer's targets of interest. For example, if Pfizer has deeply screened libraries against a class of targets (e.g. GPCRs or novel viral epitopes), its retained data could help "teach" the model nuances of those classes. The announcement suggests the model will be "tailored to Pfizer's workflows," implying integration with their internal drug development pipelines (e.g. filters for developability or compatibility with manufacturing).

This fits into a growing pattern: pharma companies using generative AI as *platform technology* that can be customized. Instead of treating AI as a black-box tool, Pfizer appears to be setting up a semi-sliding scale where Chai provides the core engine and Pfizer appends its knowledge. This may entail data-sharing under strict governance (pharma confidentiality) and possibly cloud deployment on secure infrastructure. Such arrangements raise concerns (data privacy, IP rights) but typically are structured so that the pharma retains rights to new findings (e.g. any discovery can be patented by Pfizer). Reports indicate that in MD AI partnerships, "custom models" often land proprietary control of model improvements in the sponsor's hands ([3] qz.com) ([1] www.biospace.com).

In sum, the Pfizer–Chai deal is more than a one-time software purchase; it is an ongoing collaboration. Pfizer is investing in Chai's platform and helping shape it. This dual approach allows Pfizer to benefit from Chai's generative breakthroughs while ensuring the results fit Pfizer's strategic needs. It effectively de-risks Pfizer's AI implementation: they avoid the hardest part (inventing the new algorithms) but still guide the end-product. Table 1 catalogues this and other recent pharma-AI partnerships.

Company	Partner(s)	Scope / Notes	Type of Deal	Date	Reported Value (USD)
Pfizer	Chai Discovery	AI antibody design (Chai-3, custom model)	License/Partnership	Jun 2026	*Undisclosed*
Merck	Google Cloud	Enterprise AI across R&D, manufacturing, commercial ops ([6] www.fiercebiotech.com)	Strategic Partnership	Apr 2026	Up to $1B
Novo Nordisk	OpenAI	AI integration (R&D to commercial insights) ([22] www.novonordisk.com)	Strategic Alliance	Apr 2026	*Undisclosed*
Eli Lilly	NVIDIA & InsilicoMed	AI co-innovation lab (NVIDIA $1B) and drug R&D licensing (Insilico $2.75B) ([5] www.techtarget.com)	Investments/Partnership	Jan 2026	$3.75B (total)
Alnylam	Inceptive Nucleics	Generative ML for RNAi therapeutic design	Partnership	Jun 2026	Up to $2B (Incentives)
Shionogi (Japan)	Hitachi	GenAI and data services for drug discovery	Collaboration	Jan 2025	*N/A (Press release)*
Chugai (Japan)	SB Intuitions with SoftBank	Generative AI for clinical development	Research Joint Study	Jan 2025	*N/A (Press release)*

Table 1: Selected recent partnerships between major pharmaceutical companies and AI technology providers. Deals focus on applying generative AI and related tech to drug discovery and development. Reported values are contract announcements or maximum potential.

Build vs. Buy (Partner) in Pharma AI

The Pfizer–Chai case illustrates a strategic choice: **partner with an AI specialist** versus building the technology completely in-house. In general, pharma leaders today weigh three strategic paths for AI capabilities:

- **Build (In-House Development):** Create AI solutions using internal teams and infrastructure.
- **Buy (Commercial Off-The-Shelf Solutions):** License or subscribe to existing AI tools/platforms from vendors.
- **Partner (Co-Develop / Joint Venture):** Collaborate with specialized firms or consortia to jointly develop AI.

Each path has trade-offs in cost, speed, control, and risk ([20] intuitionlabs.ai) ([21] intuitionlabs.ai). **In-house builds** offer maximum customization and data ownership: the AI can be perfectly aligned with proprietary processes and rare data. However, they require hiring scarce AI talent, curating enormous datasets, and investing potentially years before yielding results ([20] intuitionlabs.ai) ([23] eularis.com). For example, Novartis has invested heavily in an AI innovation lab and hired dedicated AI scientists, aiming for integrated in-house drug discovery platforms. But as analysts note, achieving world-class performance in a core pipeline area often takes *years and tens of millions* of investment.

By contrast, **buying turnkey solutions** yields fast deployment and lower initial cost. Off-the-shelf AI products (like cloud-based LLM APIs, analytics suites, or niche biotech platforms) can be integrated in weeks or months ([24] eularis.com) ([20] intuitionlabs.ai). They come with expert support and compliance features built-in. The downside is *vendor lock-in* (difficulty of decoupling later) and limited tailoring: most clients get a “one-size-fits-most” solution that may not address highly specific needs ([23] eularis.com) ([20] intuitionlabs.ai). For instance, many companies use GPT or other LLMs for literature analysis or customer insights simply because those tools exist, even if they only partially fit domain requirements.

Partnering is the hybrid strategy: licensing or co-developing with a specialist, as Pfizer did with Chai. In this model, the pharma sponsors or guides technology development while the partner provides deep expertise and continues to innovate. This can include equity investment, licensing deals, or joint R&D projects. The goal is to get closer to “build” in terms of customization (through collaboration and sometimes data sharing) while leveraging the partner’s foundation. Analysts like those at AI Ireland and Killer Causaly emphasize that the optimal approach is often *mix-and-match*: “Buy where the use case is generic, build where it is core to [your] competitive edge, and partner where you need specialist capability quickly” ([21] intuitionlabs.ai). In practice, companies might start with a commercial pilot (buy), then co-develop advanced features (partner), and only build fully bespoke systems for mission-critical functions. This portfolio approach spreads risk and leverages external innovation.

Below we outline key considerations and examples for each option:

- **Customization & Data Control:** Building or partnering gives more control over proprietary data. Off-the-shelf tools typically aggregate user data into shared models ([23] eularis.com).
- **Time-to-Value:** Off-the-shelf solutions can go live in weeks (ready APIs or cloud platforms) ([24] eularis.com). Building in-house can take quarters or years for complex AI systems ([24] eularis.com).
- **Cost:** Upfront, buying is cheaper (no development cost). But total cost of ownership may favor build if the solution is a long-term core asset (so you’re not paying recurring fees indefinitely) ([25] eularis.com) ([20] intuitionlabs.ai). Often, build requires significant initial capital, whereas buy spreads cost as subscription.
- **Maintenance and Upgrades:** Vendor solutions come with ongoing enhancements but depend on the vendor’s roadmap. In-house builds require in-house teams to continually update data and models (which can be onerous for data hygiene and re-training) ([23] eularis.com) ([26] intuitionlabs.ai).
- **Compliance and Risk:** In regulated pharma, buying established AI tools can ease compliance (they may include audit logging, validated workflows). Building your own agentic AI still requires rigorous validation and documentation. Partnering can share compliance burdens if the provider invests in regulated deployment.

Case Studies: This framework aligns with observed industry behavior. Pfizer chose *partner* with Chai, reflecting that antibody design is highly strategic but Chai provided unique capability. Merck’s $1B deal with Google is more of a *buy/partner hybrid*: Google provides agentic AI infrastructure and actively co-deploys engineers on site ([6] www.fiercebiotech.com). Novo Nordisk’s OpenAI alliance similarly leverages an external platform to boost internal R&D and

operations ([22] www.novonordisk.com). By contrast, companies like Novartis or Genentech are also integrating agents and LLMs in-house (with tech partnerships for which layer of AI is needed), essentially *building* core systems. Biotechs often license tools (e.g. biotech startups contracting Schrödinger or Benchling), reflecting a "buy" strategy for non-core capabilities.

Expert Perspectives: Industry analysts underscore that most pharma will end up using a hybrid approach ([21] intuitionlabs.ai) ([27] eularis.com). For example, a thorough Eularis analysis notes that building in-house confers full customization and eventual ownership—but only if firms can "ensure that they have everything they need, nothing they don't" regarding features ([23] eularis.com). Conversely, buying (commercial-off-the-shelf, or COTS, solutions) yields rapid ROI but may require workarounds or conceding competitive advantage to the vendor ([24] eularis.com). Critically, if the AI function is central to a company's unique value (e.g. novel target discovery), then build/partner becomes more attractive. If the AI use-case is generic (e.g. standard data analytics), buying an existing tool may suffice. Pfizer's choice suggests it considers generative antibody design as **core strategic capability**, but still recognizes that partnering with Chai accelerates it beyond what in-house teams could do alone.

Given these trade-offs, Table 2 summarizes the build/buy/partner spectrum:

Approach	Description	Pros	Cons	Examples
Build (In-House)	Develop AI capabilities internally (own data science teams, custom models)	Full customization to workflows; complete data ownership; untapped IP potential	Long development time; high cost and risk; requires scarce AI talent	Novartis AI Lab; Pfizer's own ML teams
Buy (COTS)	License or subscribe to existing AI platforms/tools (vendor-provided, no code)	Rapid deployment (weeks); lower initial cost; vendor support and updates; on-premise/cloud options	Less tailored; risk of vendor lock-in; shared learning (provides data to vendor) ([23] eularis.com)	Generic LLMs in R&D analytics; AWS Bio Discovery[1]
Partner/Joint Dev	Co-develop or license specialized solutions with experts (joint R&D, equity investments, or licensing)	Leverages external expertise; faster than build; more alignment than simple buy; shared IP/governance	Requires careful IP/governance agreements; integration and coordination overhead	Pfizer–Chai (antibodies); Merck–Google (AI agent & RW); GSK–Noetik (digital twins)

Table 2: Strategic options for acquiring AI capabilities in pharma (simplified). Partnerships/co-development (e.g. licensing Chai's tech) blend benefits of build and buy: gaining specialized models and support while allowing customization with internal data. ([1]AWS Bio Discovery is a recent example of a cloud platform combining multiple AI tools for antibody design).

Industry Trends and Case Studies

Pfizer's move is part of a **tectonic shift** in pharma to leverage AI at scale. Several high-profile deals in early 2026 underscore the trend:

- **Merck & Google Cloud (Apr 2026):** A multi-year collaboration where Merck invests up to $1 billion to deploy Google's agentic AI platform across drug R&D, manufacturing, and commercial operations ([6] www.fiercebiotech.com). Google's team co-works with Merck to integrate Gemini Enterprise tools for tasks like R&D workflow automation and predictive analytics ([6] www.fiercebiotech.com). This reflects Merck outsourcing heavy AI infra to a tech partner, focused on enterprise-wide gains.
- **Novo Nordisk & OpenAI (Apr 2026):** A "strategic partnership" to apply OpenAI's generative models company-wide—from drug discovery to commercial and operations ([22] www.novonordisk.com). Novo's CEO emphasized AI enabling analysis at "previously impossible" scales to uncover new therapies ([22] www.novonordisk.com). ND should handle strict data governance. Novo is also training its workforce on AI usage. This partnership is notable for using a pure tech partner (OpenAI) rather than building all in-house.
- **Eli Lilly & NVIDIA/Insilico (2026):** Lilly announced $1B co-innovation with NVIDIA for an AI research lab, plus a $2.75B licensing deal with Insilico for a portfolio of AI-discovered molecules ([5] www.techtarget.com). Lilly thus combined internal R&D investment with both tech and biotech partners. The Insilico agreement (announced May 2026) specifically gives Lilly rights to multiple preclinical assets discovered by AI.

- **AWS and Third-Party Labs:** Amazon's own "Bio Discovery" platform (April 2026) exemplifies how cloud providers are creating AI drug discovery tools that pharma can adopt. In a high-profile use case, an MSK Cancer Center researcher used AWS Bio Discovery's agent to design ~300,000 antibody variants and selected the top 100K for experimental testing – completing what normally takes a year in "just a few weeks" ([28] www.techtarget.com). This kind of service may also be "bought" by pharma or used in partnerships.

Other notable examples include collaborations like GSK's AI efforts (Noetik for digital twins), Novo's partnerships with other AI firms, and numerous internal initiatives (e.g. Siemens/Medtronic adopting AI in manufacturing). Some companies have also **acquired** AI-focused biotechs (e.g. Roche/Flatiron for data, J&J's interest in pathAI) or invested in open-source bio-AI consortia. These moves all revolve around the build-vs-partner theme: leveraging external innovation to pursue new R&D frontiers.

Empirical evidence shows these approaches pay off. For instance, *before* AWS Bio Discovery, the antibody discovery process could easily span a year with countless reagent rounds, whereas the AI-augmented process took weeks for MSK's project ([28] www.techtarget.com). Similarly, Chai's report of solving a $5M problem in hours underscores the cost and time compression possible. In purely strategic terms, deploying generative AI may reverse the famous "Eroom's law" in pharma R&D; as investor Elliot Hershberg notes, AI could fundamentally change the attrition curves by learning cumulatively from data ([29] centuryofbio.com).

However, experts caution that challenges remain. Data readiness is a common bottleneck: integrating silos (clinical, preclinical, literature) is tedious and often underlies AI project delays ([30] intuitionlabs.ai) ([20] intuitionlabs.ai). Human validation is still required – AI designs currently need yeast display or animal tests to confirm efficacy (as in Chai-2 and RFdiffusion studies) ([12] pmc.ncbi.nlm.nih.gov). Moreover, not all companies have equal data assets; this may widen gaps between well-resourced pharmas (like Pfizer) and smaller firms. Finally, AI models must navigate regulatory scrutiny (FDA draft guidance on AI tools was issued in 2025), making transparency and reliability paramount.

Nevertheless, the momentum is undeniable. Within weeks of Pfizer's announcement, industry news sites were already heralding the era of generative antibody design as "coming soon in pharma" ([31] www.techtarget.com). Pharma CEOs have publicly acknowledged AI's central role – for example, Pfizer's Albert Bourla commented on using technology to bolster pipeline productivity alongside growth ([13] qz.com). As one analyst summary put it, major drugmakers are rapidly shifting "from research breakthroughs into deployment across real discovery workflows" ([32] www.biospace.com). The practical evidence is mounting: companies integrating AI report faster candidate progression, earlier target ID, and new target classes explored. For example, Chai-2 not only found antibodies to well-known targets, but also generated binders against **miniproteins** and entirely novel epitopes (colab.ws) ([12] pmc.ncbi.nlm.nih.gov), hinting at future capabilities beyond traditional biologics.

Discussion: Implications and Future Directions

The Pfizer–Chai alliance, and its contemporaries, signal that generative AI has "arrived" in pharma. The **implications** for drug development are profound:

- **Acceleration of Discovery:** As shown by Chai and AWS results, AI can collapse timelines from months/years into weeks. This will pressure all programs to adopt AI or risk falling behind. In antibody discovery, it could mean the difference between a lead in 6 months versus 2 years. More broadly, generative AI for small molecules (as with other startups like Exscientia) is similarly poised to expedite hit-finding.
- **Expanded Target Space:** AI design can tackle targets once considered "undruggable." Chai's platform explicitly aims at "targets traditional methods struggled to reach" ([19] www.biospace.com). For example, in silico models could design binders to buried or complex epitopes, or multi-drug formats (bi/trispecifics) that ordinary methods rarely produce. Over time, this may yield therapies for diseases that have eluded conventional biotech, from complex cancers to novel viral threats.

- **New R&D Paradigms:** The pharma R&D model may shift to a more **integrated computational-experimental loop**. Companies will rely on in silico design wheels driven by AI, with relatively light experimental screening to filter top designs. Iterative feedback (lab data feeding AI retraining) can become continuous. This agile approach echoes the "DevOps" model in software, applied to biology. Tools like AWS Bio Discovery epitomize this "lab-in-the-loop" vision ([33] www.techtarget.com).
- **Economic Impact:** If AI can systematically reduce attrition and speed time-to-clinic, it could substantially improve the economics of drug development. Conservative estimates suggest AI could add tens of billions in value to pharma R&D efficiency ([34] intuitionlabs.ai). The high investment levels confirm industry expectations of huge returns.
- **Regulatory Evolution:** Regulators will need to adapt. AI-generated drug candidates challenge existing frameworks (e.g. how to document AI's role in target validation, how to ensure reproducibility of design). We note the nascent FDA guidance and impending EU AI Act, which will classify some AI tools as "high-risk" in healthcare. Companies will need robust validation pipelines and explainability for their AI components. This could actually favor partnerships: specialized AI vendors may develop compliant infrastructure faster than each pharma reinventing it.
- **Talent and Organization:** The workforce must evolve to exploit these tools. Drug scientists will need AI literacy. Organizationally, cross-functional teams (computational biologists, chemists, data scientists, domain experts) must collaborate seamlessly. This cultural shift is non-trivial and will be a key bottleneck.
- **Strategic Positioning:** How a pharma navigates the build/partner choice could become a competitive differentiator. Companies that effectively integrate external innovation (like Pfizer with Chai) may outpace those clinging to slower internal builds. Yet, reliance on partners might also cede some unique insights. The "optimal strategy is hybrid" guidance ([21] intuitionlabs.ai) will likely be tested in practice: firms may spin up pilot AI units (buy), then selectively co-develop (partner) for strategic projects, while reserving in-house teams for proprietary platform elements.

Finally, we note the historical context: just as the industrial revolution shifted manufacturing paradigms, the AI revolution is shifting drug discovery from trial-and-error to algorithm-driven design. Chai Discovery's rise (backed by top AI investors) underscores this sea change. We anticipate continued convergence: LLMs (like GPT-based tools) will become integrated with molecular AIs, enabling new kinds of multimodal discovery agents. The "zero-shot" concept itself may evolve: future models might even predict efficacy or toxicity profiles in harmony with binding. And as more case studies accumulate, best practices for AI adoption will emerge (e.g. standardized benchmarks, regulatory pathways).

Conclusion

Pfizer's licensing of Chai-3 and the related custom AI model marks a **milestone** in the integration of generative AI into pharmaceutical R&D. It validates that companies no longer view AI as a fringe experiment but as an essential component of the drug discovery toolkit. The Chai story illustrates both the opportunities of cutting-edge AI – turning an intractable problem into a solved one in hours – and the strategic choices facing pharma executives. By partnering with Chai, Pfizer has effectively bought into the future of biologics design. Whether they will fully build similar capabilities internally is an open question; for now, partnership has allowed them to leap forward.

Looking ahead, we foresee a landscape where build, buy, and partner mix fluidly. Some firms will develop proprietary AI foundations, while others will rely on open-source models and cloud platforms, and most will do a combination. Data sharing agreements and ethical guidelines will need to keep pace. Yet the direction is clear: generative models like Chai-3 will become a standard R&D engine. As AI discovery tools become more sophisticated, we expect new classes of therapeutics to emerge and a gradual reshaping of pharmaceutical pipelines.

In sum, Pfizer–Chai is a case study in modern pharma strategy: leveraging state-of-the-art science through targeted partnership. It provides a blueprint – the "Build-vs-Partner AI playbook" – for others to follow. Through extensive citations and analysis, this report has shown that Chai's generative antibody design platform exemplifies the technological promise of AI and the pragmatic considerations of deployment in industry. As the field advances, the companies that successfully integrate AI with biological insight will likely set the pace for the next generation of medicines.

References: All claims above are supported by cited sources. Key facts about Chai-2/Chai-3 performance come from Chai Discovery's publications ([4] finance.yahoo.com) (colab.ws). Details of Pfizer's license and Chai-3 capabilities are from

Business Wire and Quartz reports ([1] www.biospace.com) ([2] www.businesswire.com). Industry partnerships and strategic analysis are drawn from news articles and expert analyses ([6] www.fiercebiotech.com) ([22] www.novonordisk.com) ([20] intuitionlabs.ai). Additional context on AI in pharma is cited from reviews and press ([7] pmc.ncbi.nlm.nih.gov) ([28] www.techtarget.com) ([21] intuitionlabs.ai). Tables and analysis synthesize these findings.

External Sources

[1] https://www.biospace.com/press-releases/chai-discovery-announces-license-agreement-with-pfizer-to-accelerate-drug-discovery-with-ai#:~:The%2...

[2] https://www.businesswire.com/news/home/20260602498831/en/Chai-Discovery-Announces-License-Agreement-with-Pfizer-to-Accelerate-Drug-Discovery-with-AI#:~:secur...

[3] https://qz.com/pfizer-chai-discovery-ai-drug-discovery-license-060526?.tsrc=rss#:~:The%2...

[4] https://finance.yahoo.com/news/chai-discovery-unveils-chai-2-100000398.html#:~:annou...

[5] https://www.techtarget.com/pharmalifesciences/news/366641915/AWS-launches-AI-driven-tool-to-speed-up-early-stage-antibody-discovery#:~:also%...

[6] https://www.fiercebiotech.com/ai-and-machine-learning/merck-goes-google-ai-push-1b-enterprise-deal#:~:As%20...

[7] https://pmc.ncbi.nlm.nih.gov/articles/PMC11254834/#:~:for%2...

[8] https://finance.yahoo.com/news/chai-discovery-unveils-chai-2-100000398.html#:~:Exten...

[9] https://finance.yahoo.com/news/chai-discovery-unveils-chai-2-100000398.html#:~:Chai,...

[10] https://finance.yahoo.com/news/chai-discovery-unveils-chai-2-100000398.html#:~:%22Ch...

[11] https://pmc.ncbi.nlm.nih.gov/articles/PMC11254834/#:~:Compa...

[12] https://pmc.ncbi.nlm.nih.gov/articles/PMC12727541/#:~:The%2...

[13] https://qz.com/pfizer-chai-discovery-ai-drug-discovery-license-060526?.tsrc=rss#:~:The%2...

[14] https://pmc.ncbi.nlm.nih.gov/articles/PMC11254834/#:~:Antib...

[15] https://pmc.ncbi.nlm.nih.gov/articles/PMC11254834/#:~:Recen...

[16] https://finance.yahoo.com/news/chai-discovery-unveils-chai-2-100000398.html#:~:Story...

[17] https://finance.yahoo.com/news/chai-discovery-unveils-chai-2-100000398.html#:~:,medi...

[18] https://www.biospace.com/press-releases/chai-discovery-announces-license-agreement-with-pfizer-to-accelerate-drug-discovery-with-ai#:~:This%...

[19] https://www.biospace.com/press-releases/chai-discovery-announces-license-agreement-with-pfizer-to-accelerate-drug-discovery-with-ai#:~:%E2%8...

[20] https://intuitionlabs.ai/articles/agentic-ai-pharma-build-buy-framework#:~:how%2...

[21] https://intuitionlabs.ai/articles/agentic-ai-pharma-build-buy-framework#:~:As%20...

[22] https://www.novonordisk.com/content/nncorp/global/en/news-and-media/news-and-ir-materials/news-details.html?id=916532#:~:The%2...

[23] https://eularis.com/build-versus-buy-for-pharmaceutical-ai-should-i-build-or-should-i-buy/#:~:Custo...

[24] https://eularis.com/build-versus-buy-for-pharmaceutical-ai-should-i-build-or-should-i-buy/#:~:Quick...

[25] https://eularis.com/build-versus-buy-for-pharmaceutical-ai-should-i-build-or-should-i-buy/#:~:,of%2...

[26] https://intuitionlabs.ai/articles/agentic-ai-pharma-build-buy-framework#:~:its%2...

[27] https://eularis.com/build-versus-buy-for-pharmaceutical-ai-should-i-build-or-should-i-buy/#:~:,Busi...

[28] https://www.techtarget.com/pharmalifesciences/news/366641915/AWS-launches-AI-driven-tool-to-speed-up-early-stage-antibody-discovery#:~:Nai,w...

[29] https://centuryofbio.com/p/eroom#:~:AI%20...

[30] https://intuitionlabs.ai/articles/agentic-ai-pharma-build-buy-framework#:~:,26...

[31] https://www.techtarget.com/pharmalifesciences/news/366641915/AWS-launches-AI-driven-tool-to-speed-up-early-stage-antibody-discovery#:~:Accor...

[32] https://www.biospace.com/press-releases/chai-discovery-announces-license-agreement-with-pfizer-to-accelerate-drug-discovery-with-ai#:~:Chai%...

[33] https://www.techtarget.com/pharmalifesciences/news/366641915/AWS-launches-AI-driven-tool-to-speed-up-early-stage-antibody-discovery#:~:With%...

[34] https://intuitionlabs.ai/articles/agentic-ai-pharma-build-buy-framework#:~:%28%5...

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.