IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

NVIDIA-Eli Lilly AI Lab: Drug Discovery Compute Strategy

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

Published by IntuitionLabs · 2026-04-18 · ai drug discovery · nvidia eli lilly · pharma infrastructure · computational biology · dgx superpod · foundation models · bionemo · exascale computing

Original article: https://intuitionlabs.ai/articles/nvidia-eli-lilly-ai-drug-discovery-infrastructure



In this article

1. Executive Summary
2. Introduction and Background
3. NVIDIA–Eli Lilly Co-Innovation Lab: Vision and Structure
4. Joint Objectives and Scope
5. Infrastructure: AI “Factory” and Robotics
6. Data, Models, and Software Stack
7. 5-Year Compute Strategy
8. Heritage and Case Studies
9. Infrastructure and Compute Details
10. Data Generation and AI Models
11. Compute Performance and Statistics
12. Case Studies and Use-Cases
13. Implications, Challenges, and Future Directions
14. Table: Key Technologies in the NVIDIA–Lilly AI Lab
15. Data Analysis and References
16. Discussion and Future Directions

Executive Summary

- **NVIDIA–Eli Lilly AI Co-Innovation Lab (2026):** In January 2026, NVIDIA and Eli Lilly announced a landmark partnership to create a dedicated “AI co-innovation lab” focused on drug discovery (investor.lilly.com) (www.theregister.com). The lab—based in the San Francisco Bay Area—is backed by *up to $1 billion* in joint investment over five years in compute infrastructure, talent, and research (investor.lilly.com). Its mission is to **re-invent the drug discovery process** using advanced AI and compute, bringing together Lilly’s biomedical domain expertise with NVIDIA’s AI hardware/software leadership (investor.lilly.com).
- **Key Technologies and Infrastructure:** The lab will build on NVIDIA’s state-of-the-art platforms, including the **BioNeMo** AI framework for biomedicine and the upcoming **Vera Rubin** hardware architecture (investor.lilly.com) (www.noah.bio). Central to the plan is an exascale-class **AI “factory” supercomputer** – a DGX SuperPOD containing 1,016 B300 Blackwell Ultra GPUs (NVIDIA operational report) (investor.lilly.com) – capable of training foundation models on Lilly’s vast proprietary datasets. Robotics and automation (“physical AI”) will link Lilly’s **wet labs** and **dry labs** in a continuous 24/7 closed-loop, enabling scientist-in-the-loop experimentation (investor.lilly.com) (nvidianews.nvidia.com). Digital twin technology (using NVIDIA Omniverse and RTX servers) will simulate Lilly’s manufacturing and supply chains before real-world deployment (nvidianews.nvidia.com).
- **5-Year Compute Strategy:** The announced co-innovation lab infrastructure is to be built on NVIDIA BioNeMo and the NVIDIA Vera Rubin architecture (NVIDIA announcement). This is architectural intent; it does not establish a Lilly deployment schedule, exponential throughput gains or a measured tenfold

reduction in drug-screening costs.

- **AI-Driven Discovery Workflow:** The lab aims to transform R&D from a traditional “artisanal” process into an *engineering discipline*. By training next-generation **foundation and frontier models** on Lilly’s decades of omics and experimental data, scientists can explore vast molecular and chemical spaces in silico (investor.lilly.com) (www.noah.bio). In practical terms, the AI models will generate and screen millions of candidate molecules, optimize lead compounds, and even propose novel biologics with enhanced binding profiles (www.noah.bio) (blogs.nvidia.com). These predictions are rapidly validated by automated chemistry and biology labs, closing the “dry-to-wet” loop. In this way, Lilly and NVIDIA expect to overcome the industry’s worsening R&D productivity (the so-called “Eroom’s Law”) and dramatically accelerate the pace of new drug leads (nvidianews.nvidia.com) (www.noah.bio).
- **Benchmarking and Ecosystem Context:** The Lilly–NVIDIA lab will join a small but growing set of large-scale AI/HPC initiatives in life sciences. LillyPod contains 1,016 B300 Blackwell Ultra GPUs (NVIDIA operational report). Raw accelerator counts alone do not establish equivalent application performance across different systems. (See **Table 1**.) Similarly, tech giants like Google and Amazon are deploying AI platforms for drug discovery: Google DeepMind’s AlphaFold models (cited with a 2024 Nobel Prize) predict protein structures at atomic accuracy (time.com) (www.lemonde.fr), and AWS’s **BioDiscovery** service offers 40+ specialized AI models to triage antibody candidates (www.techradar.com). The Lilly lab anchors NVIDIA’s broader healthcare strategy (including Thermo Fisher collaboration and open AI coalitions) and positions Lilly among companies (e.g. Novartis, AstraZeneca, Bayer) racing to integrate AI into biology.
- **Implications and Future Directions:** This co-innovation lab represents a strategic pivot: *vast compute power + AI = reimagining pharma R&D*. If successful, it could shorten discovery timelines, bring more candidates to the clinic, and maintain Lilly’s competitiveness in an era where “AI-first” biotech firms (Insilico, Insitro, etc.) are emerging. It also exemplifies the “ agentic” future of labs—virtualized robots working 24/7—and foreshadows a new digital infrastructure for medicine. However, challenges remain: validating AI predictions, managing experimental errors, and ensuring models generalize. On balance, experts agree this marks a major step in turning long-standing drug discovery “art” into a data-driven engineering discipline (www.prnewswire.com) (www.lemonde.fr).

Introduction and Background

The Drug Discovery Challenge: Discovering a new therapeutic molecule has long been *incredibly costly and slow*. Industry studies estimate that bringing a novel drug to market typically takes over a **decade or more** of R&D (apnews.com) and incurs development costs that vary substantially by methodology. CBO’s 2021 review reports estimates ranging from **less than $1 billion to more than $2 billion**, including capital costs and spending on drugs that fail to reach the market (CBO review). These immense costs arise from complex biology, high failure rates in trials, and labor-intensive experimentation. In recent decades, pharmaceutical productivity has counterintuitively **declined** (“Eroom’s Law” – the inverse of Moore’s Law) even as basic science and data have grown (nvidianews.nvidia.com) (apnews.com). The COVID era and genomics

revolution only magnify the need: researchers now can generate petabytes of biological data (gene sequences, proteomics, high-throughput screening, etc.), far surpassing the capacity of traditional lab workflows to fully exploit.

AI and HPC in Pharma: Over the last several years, the industry has recognized artificial intelligence as a key to breaking this impasse. Machine learning models can sift massive datasets (experimental assays, molecular databases, patient biomarkers) to spot patterns unattainable by humans alone. Early successes—such as **DeepMind's AlphaFold**, a neural network that predicts protein 3D structures with near-experimental accuracy (www.lemonde.fr)—have galvanized the field. In late 2024, the Nobel Prize in Chemistry was partly awarded to AI-driven structural biology (recognizing AlphaFold's impact (www.lemonde.fr)). Meanwhile, startups like **Insitro** and **Insilico** use generative AI to propose candidate molecules, and companies (AstraZeneca, GSK, Novartis, etc.) have announced AI partners. For example, pharmaceutical company Eli Lilly has already struck a multi-hundred-million-dollar deal with Hong Kong's Insilico Medicine on AI-driven drug design (omniekonomi.se), and agreements with San Francisco's Insitro (apnews.com).

Alongside AI algorithms, ultra-fast computing hardware has become equally critical. Modern deep learning often requires training **"foundation models"** (large neural networks) on hundreds of billions of data points, demanding *exa*-scale compute (10^18 operations). NVIDIA, a leader in GPUs and supercomputers, has seized on life sciences as a key growth sector. The company has already collaborated with national labs (Oak Ridge, Los Alamos, Europe's Jülich) and institutional partners (Riken, etc.) to develop exascale AI supercomputers (www.theregister.com) (www.noah.bio). Now, the partnership with Eli Lilly signals a move into industrial drug discovery: it brings together **one of the deepest GPU compute stacks in the world with one of the largest pharmaceutical R&D pipelines**. As Lilly CEO Dave Ricks put it, combining Lilly's vast datasets and biological expertise with NVIDIA's computational power "could reinvent drug discovery" (investor.lilly.com).

NVIDIA's Strategic Push in Healthcare: NVIDIA has been explicitly building an AI platform for biomedicine over the last few years. In 2023–2024, it launched the **BioNeMo** (Bio-Pharmaceutical Neural Modeling) software suite and released pretrained **foundation models** for genomics, proteomics, and small molecules (blogs.nvidia.com). Its Clara and Omniverse tools target medical imaging and visualization. The Lilly lab will leverage this existing ecosystem. NVIDIA CEO Jensen Huang has publicly forecast that "drug research will shift from traditional labs to AI-driven platforms" (www.axios.com), citing Lilly as an eager partner. (In January 2026, at Davos, Huang noted Lilly is building "a supercomputer capable of developing research models" for medicine (www.axios.com).)

Against this backdrop, the **NVIDIA–Eli Lilly AI Co-Innovation Lab** was launched. It is described as *"first-of-its-kind"*, combining a biomedical giant's knowledge with cutting-edge computational infrastructure (investor.lilly.com). Between Lilly's historical strengths (first $1 trillion Pharma company (www.axios.com), 150 years of R&D (investor.lilly.com), a business focused on innovative medicines (Lilly 2025 Form 10-K), and a recent $6B biopharma plant investment (www.axios.com)) and NVIDIA's leadership in AI (first $5 trillion valuation company, architecture IP, GPU supply), the collaboration aims to tackle head-on the "hardest problems" in pharmaceutical discovery (investor.lilly.com) (www.theregister.com).

NVIDIA–Eli Lilly Co-Innovation Lab: Vision and Structure

Joint Objectives and Scope

The co-innovation lab was **announced on January 12, 2026** at the JPMorgan Healthcare Conference; the announcement expected work to begin in South San Francisco early that year (NVIDIA announcement). It is co-located in the San Francisco Bay Area, intentionally mixing Lilly's biologists, chemists and manufacturing experts with NVIDIA's AI scientists and engineers (investor.lilly.com). By co-housing teams, the companies seek a fast feedback loop: hypotheses and models developed in close proximity can be rapidly tested and iterated.

Key **goals** of the lab include:

- Developing *next-generation AI models* for biology and chemistry (both **foundation** models and specialized **frontier** models) that can predict candidate molecules, drug-target interactions, and biological effects (investor.lilly.com).
- Constructing a *continuous learning system*, where experimental "wet lab" data flows directly into AI training and vice versa (investor.lilly.com). For example, a prediction of a promising molecule is quickly synthesized and tested, then the result labels further training data, in a 24/7 cycle.
- Pioneering **agentic AI** and robotics: autonomous scientific "agents" will design and execute experiments. NVIDIA CEO Jensen Huang said AI's most profound impact would be in life sciences (NVIDIA announcement).
- Applying AI beyond basic discovery: the partnership will explore AI in later stages too—clinical development, manufacturing optimization, and even commercial operations (nvidianews.nvidia.com). In particular, they will use **digital twin** technologies (via NVIDIA Omniverse and RTX hardware) to simulate Lilly's entire manufacturing lines and supply chains, enabling virtual stress-testing of processes before costly real-world changes (nvidianews.nvidia.com).

The official press release summarizes: *"Building a Continuous Learning System for Drug Discovery"*, connecting Lilly's **agentic wet labs** with **computational dry labs** (investor.lilly.com). This "scientist-in-the-loop" paradigm is envisioned to break traditional R&D bottlenecks. By the same token, Lilly's Chief AI Officer Thomas Fuchs stated that Lilly is moving "from using AI as a tool to embracing it as a scientific collaborator… embedding intelligence into every layer of our workflows" (www.prnewswire.com). In short, the lab's blueprint is to treat the drug pipeline itself as an **engineering system**: data + compute + automation, iterated continuously to produce breakthroughs neither company could achieve alone (investor.lilly.com).

Infrastructure: AI "Factory" and Robotics

LillyPod is Lilly's NVIDIA DGX SuperPOD with DGX B300 systems. It contains **1,016 B300 Blackwell Ultra GPUs**, rather than more than 1,000 DGX servers (NVIDIA operational report). A DGX B300 server contains eight B300 GPUs (NVIDIA DGX B300 guide).

This AI supercomputer feeds into Lilly's **federated research platform (TuneLab)**, which will host the trained models and workflows. The October 2025 news noted that scientists at Lilly could train models on *millions of experiments* simultaneously, vastly expanding discovery scope (investor.lilly.com). By way of example from that release: Lilly expects to make its proprietary AI models available on TuneLab to the broader biopharma ecosystem, alongside NVIDIA's open models (e.g. Clara) (investor.lilly.com). Moreover, the high-throughput computing will not be limited to chemistry: computing power is also earmarked for advanced applications like medical imaging AI (leveraging NVIDIA Clara) and internal "AI agents" that plan and coordinate research (nvidianews.nvidia.com) (investor.lilly.com).

Beyond the static compute cluster, the lab will embed **robotics and physical AI**. Lilly is already investing in automated lab equipment and "robot scientists" for synthesis and screening. Under the new partnership, NVIDIA's robotics interfaces (via Omniverse and other toolkits) will integrate these. This setup will allow models to generate a hypothesis (e.g. "test molecule X") and then dispatch robots to carry out synthesis and bioassays, with results automatically fed back. As the press materials state, NVIDIA and Lilly *"will pioneer robotics and physical AI to accelerate and scale medicine discovery"* (nvidianews.nvidia.com) (investor.lilly.com). In practice, this means developing a **dry-wet loop**: AI-designed experiments take place around the clock in Lilly labs, far exceeding human throughput. Lilly's Chief AI Officer summarized: embedding AI at every layer lets their R&D enterprise "learn, adapt and improve with every data point (www.prnewswire.com)," moving beyond merely speeding up tasks to actually probing biology at scale.

Data, Models, and Software Stack

Lilly brings to the table **decades of proprietary biomedical data**: patient genetics, high-throughput screening results, longitudinal clinical data, chemical libraries, imaging data, etc. These large, high-quality datasets will be the fuel for NVIDIA's software. The co-innovation lab infrastructure is explicitly built on the **NVIDIA BioNeMo platform**, which bundles pre-trained foundation models and tools for drug discovery (nvidianews.nvidia.com). BioNeMo provides libraries (e.g. nvMolKit for cheminformatics) and "recipes" for common tasks like molecular generation, protein folding proxies, and sequence analysis (techlife.blog). Under the lab, researchers will fine-tune these BioNeMo models on Lilly's data, creating what technical documentation calls "specialized foundation models on [Lilly's] decades of longitudinal and multi-omic data" (www.noah.bio).

Practically, this means the lab will train **biology foundation models** (large-scale neural nets) that capture relationships between molecular structure and biological activity. For example, a model might take as input a protein target and a candidate drug and predict binding affinity or off-target effects. Another model might analyze genetic sequences to identify indicators of disease. NVIDIA has already released BioNeMo models capable of tasks like DNA/RNA sequence analysis and protein shape prediction (blogs.nvidia.com), and Lilly's lab will extend these with Lilly's vast internal datasets. The joint announcement describes a framework in which experiments, data generation and AI model development continuously inform and improve one another (Lilly–NVIDIA announcement). This pipeline essentially automates the **design-build-test-learn (DBTL)** cycle at extreme scale.

A crucial aspect is that the models will be **“multimodal” and agentic**. Beyond pure molecular graphs, NVIDIA and Lilly plan to integrate chemical, genomic, image, and text data. For sequence analysis, the ESM-2nv protein language model provides numerical representations of each amino acid in a protein sequence (NVIDIA model documentation). NVIDIA has hinted at integrating “scientific AI agents” that not only predict outcomes but also suggest experiments (nvidianews.nvidia.com). In other words, the AI could propose next steps (akin to an AI research assistant) and coordinate with lab robotics. This approach draws on research showing that transformer-based “foundation” models can indeed learn general principles of biology, as seen in Google’s work: AlphaFold 3 (released May 2024) is an example of a foundation biology model that predicts protein interactions (time.com). In summary, the software stack will include BioNeMo frameworks plus NVIDIA’s other tools (Clara for imaging, Omniverse for 3D simulations) under an NVIDIA AI orchestration, all tuned to Lilly’s use cases.

5-Year Compute Strategy

The announced co-innovation lab infrastructure is to be built on **NVIDIA BioNeMo and Vera Rubin** (NVIDIA announcement). This architectural intent should not be treated as a contractual guarantee of successive hardware upgrades or a published deployment schedule.

LillyPod brings together proprietary data and advanced AI models (NVIDIA operational report).

Besides raw hardware, the 5-year strategy encompasses software and platform evolution. NVIDIA plans to continuously advance the BioNeMo toolkit (e.g. new pretrained models, optimized libraries) while Lilly will invest in data pipelines (curation, labeling, federated access). On the computing side, the early reliance on an on-prem SuperPOD might be complemented by cloud bursting (e.g. via NVIDIA’s DGX Cloud or equivalently large NVIDIA Cloud instances) during peak training periods. Indeed, NVIDIA’s broader ecosystem initiatives (Nemotron coalition with open frontier models , DGX Cloud services, etc.) suggest Lilly could tap community resources too.

Another component is **energy and sustainability**. Lilly has emphasized carbon-neutral commitments; the Oct 2025 PR notes the supercomputer will run on 100% renewable power within Lilly’s campuses and use efficient liquid cooling (www.prnewswire.com). This is a non-negligible part of the compute strategy: continually scaling supercomputing needs to balance performance with energy usage and cost. By planning for green power and efficient designs, the lab aims to mitigate the high energy footprint of AI training. (Notably, Isambard-AI is already one of the world’s most energy-efficient at rank #4 on the Green500 (www.techradar.com).)

Finally, the strategy involves talent and collaboration. Lilly describes its work with NVIDIA as uniting compute, specialized talent and data at scale (NVIDIA BioNeMo announcement). Conceptually, the “compute strategy” covers **both hardware and people** to use it.

Heritage and Case Studies

To appreciate the significance of the NVIDIA–Lilly lab, it helps to consider parallel case studies where AI and compute are being applied in pharma. These illustrate both the promise and the challenges of the approach:

- **Google DeepMind – AlphaFold (Protein Folding):** In 2020, DeepMind released AlphaFold 2, a deep learning system that predicts 3D protein structure from amino acid sequences. By mid-2024, DeepMind announced **AlphaFold 3** – now capable of predicting interactions between proteins, DNA/RNA, and small-molecule drugs (time.com). The AlphaFold Protein Structure Database makes predicted protein structures freely accessible to researchers (AlphaFold DB). This AI breakthrough **earned a Nobel Prize in Chemistry in 2024** (the leaders highly cited that everyone in protein research now uses AlphaFold predictions (www.lemonde.fr)). For drug discovery, AlphaFold accelerates structure-based design: medicinal chemists no longer need to crystallize proteins (a slow process) to get formulations for docking; AI gives them a head start in minutes (www.lemonde.fr). The NVIDIA–Lilly initiative takes inspiration from this: instead of just proteins, their generative biology models will handle entire pathways and chemistries. However, whereas AlphaFold's code is publicly open, Lilly's models will be proprietary and likely larger in scale. This shows how GPU-powered AI can revolutionize biology, but also that broad adoption (via open science) was key. The co-innovation lab indicates Lilly wants its own advantage rather than rely on open models alone.
- **AWS BioDiscovery – Lab-in-the-Loop:** In 2026, Amazon Web Services launched **Amazon BioDiscovery**, a cloud platform with over 40 AI models specialized for antibody design and molecular screening (www.techradar.com). A key feature is the "lab-in-the-loop" workflow: AWS provides AI agents that suggest models/parameters, filters millions of candidates, and then directly interfaces with partner labs for synthesis and testing, eliminating slow handovers (www.techradar.com). BioDiscovery emphasizes democratization ("every researcher"), but it points to a trend: making AI pipelines accessible end-to-end. Nvidia and Lilly's lab goes much deeper (own hardware, top talent), but conceptually it mirrors this integration of AI suggestions with real wet lab execution (www.techradar.com). Moreover, AWS's tools underscored the need for foundational models tailored to biology (cloud catalogs of models for antibodies) – something NVIDIA's BioNeMo platform similarly aims to build. The experience of AWS hints at challenges: ease-of-use, model explainability, and costs. Lilly's in-house approach avoids some AWS hurdles (network latency, data privacy on cloud) but must solve them in its own infrastructure.
- **Insitro and Insilico Ventures:** Both Insitro (US) and Insilico Medicine (HK) are biotech startups using generative AI for small molecules and proteins. Insitro (founded 2018) has multi-year deals with large pharmas including Eli Lilly and Bristol Myers Squibb to use ML on genetic and phenotypic data (apnews.com). Insitro's founder and CEO is **Daphne Koller** (Insitro). Insilico Medicine's March 29, 2026 announcement describes eligibility for **$115 million upfront**, with development, regulatory and commercial milestones that could bring total deal value to **approximately $2.75 billion**, plus tiered royalties on future sales (exchange filing). These deals indicate Lilly's broad AI strategy: the Nvidia partnership is not her sole focus; they are also tapping multiple external AI partners. The case study here is that of a large pharmaceutical adopting both internal HPC (NVIDIA lab) and external startup collaborations to cover different angles. In practice, Insitro's work involves building predictive models for targets and phenotypes using high-dimensional data – much as the Lilly lab intends to do on a larger scale. The risk noted by experts: such models often perform well in "in silico" benchmarks, but translating to actual new drugs is unproven. The NVIDIA–Lilly lab can be seen as a bolt-on to Insitro-type approaches – providing more compute and integration into Lilly's pipelines.

- **Nvidia's Broader AI Coalitions:** NVIDIA's **Nemotron Coalition** brings together AI developers to collaboratively develop a shared model (NVIDIA coalition announcement). They also paired with Thermo Fisher Scientific to create lab infrastructure ("fundamental AI infrastructure of the lab") using DGX Spark devices (www.stocksfoundry.com). These moves show NVIDIA's play: they see biomedical research as a vast emerging user of AI/HPC. The Lilly lab thus fits into a larger trend of vendor–researcher co-development. In these coalitions, NVIDIA provides hardware and core AI tech, while domain experts (like Lilly) supply knowledge and data. From one perspective, this is analogous to how NVIDIA collaborates with national labs (e.g. Riken's FugakuNEXT) or automotive (NVIDIA DRIVE for autonomous vehicles).
- **Pharma Industry Trends:** More broadly, other big pharmas are ramping up AI. At industry summits (e.g. Axios Biotech BFD, Nov 2025), leaders from AstraZeneca, GSK, and Bayer discussed how AI is **streamlining drug R&D** (www.axios.com). Bayer's pipeline head noted that AI improved their ability to screen "gene-driven diseases" (www.axios.com). Merck's CFO has publicly mentioned Google Cloud AI collaborations. Even Intel and IBM have pitched AI systems for molecular modeling (e.g. IBM Watson). In that context, Lilly's NVIDIA deal is one of the largest yet – a $1B lab dwarfs typical point collaborations. It also represents a marriage of hardware and pharma. Industry analysts see this as emblematic of the new "digital biology" era: integrated AI/hardware companies partnering with life science firms to fundamentally change the R&D process.

These examples show two key points: (1) AI-driven biotech partnerships are proliferating; and (2) integrating massive compute is considered essential. Lilly says combining its scientific expertise with NVIDIA's AI computing can enhance how it works (Lilly's account). The NVIDIA lab can be seen as a **post-graduation** of past efforts: from small ML projects to a full-scale AI factory. Whether such labs succeed in delivering new therapies faster remains to be seen, but the direction of travel – hardware-accelerated drug design – is clear.

Infrastructure and Compute Details

The NVIDIA–Lilly lab represents one of the most advanced **drug discovery compute infrastructures** announced to date. Key components include:

- **DGX SuperPOD (Lilly AI Factory):** LillyPod contains **1,016 B300 Blackwell Ultra GPUs** (NVIDIA operational report). Each DGX B300 server contains eight B300 GPUs (NVIDIA system guide). Table 1 lists system configurations without ranking their performance.
- **Networking and Storage:** NVIDIA documents **Spectrum-X Ethernet networking** for LillyPod (NVIDIA operational report).
- **BioNeMo Platform:** The software platform is NVIDIA's BioNeMo. This includes pre-trained *foundation models* (for DNA/RNA, proteins, etc.), data-processing libraries (e.g. nvMolKit accelerating cheminformatics) (techlife.blog), and cloud services (BioNeMo Service) for pipeline management. Lilly will use BioNeMo "recipes" as starting points to fine-tune models on its own data (techlife.blog). Over time, Lilly's team will develop new models for specific tasks (e.g. small-molecule property prediction).

- **Next-Gen Accelerators (Vera Rubin):** Vera Rubin combines **Vera CPUs and Rubin GPUs** (NVIDIA technical overview). NVIDIA's January 5, 2026 launch announcement scheduled partner availability for the second half of 2026 (NVIDIA announcement).
- **Digital Twins & Omniverse:** The lab will deploy NVIDIA Omniverse libraries and RTX PRO visualization servers to create *virtual copies* of Lilly's real-world labs and factories (nvidianews.nvidia.com). These digital twins will allow engineers and managers to simulate equipment layouts, test the impact of new processes, and train AI control systems before physical changes. For instance, before retooling a production line for increased capacity, Lilly can run stress tests in Omniverse to find bottlenecks. This use of GPUs for simulation and optimization is akin to what automotive companies do for autonomous vehicle design, and what aerospace firms do for simulations; here it's applied to biomanufacturing.
- **Cloud and Edge Integration:** Although the primary compute is on-premises, the plan includes hybrid cloud. NVIDIA's DGX Cloud (a cloud service) and a software-defined fleet could be used for scaling during peak periods. Lilly has existing commitments to major cloud providers; they may extend these to NVIDIA's branded cloud instances when needed. The multi-year plan likely includes staying flexible about cloud vs local: initial model training in-house, then potentially inference or multi-site collaboration in hybrid mode (especially for collaboration with academic partners or consortiums).

Table 1. *Comparison of major AI/HPC systems relevant to drug discovery (through 2026).* （GPUs counts and flops are approximate and based on public reports.）

Supercomputer	Location / Owner	GPUs / Compute configuration	Purpose / Notes
LillyPod (February 2026 inauguration)	Eli Lilly (Indianapolis lab)	1,016 B300 Blackwell Ultra GPUs (NVIDIA operational report)	Pharmaceutical R&D; training large biomedical models (investor.lilly.com)
Isambard-AI (Jul 2025)	Bristol, UK (BriCS)	5,448 NVIDIA GH200 GPUs (www.techradar.com)	General HPC (rank #11 Top500); 4th Green500 (energy-efficient) (www.techradar.com)
JUPITER (Sept 2025)	Forschungszentrum Jülich, Germany (EuroHPC)	24,000 NVIDIA GH200 GPUs (www.pcgamer.com)	Exascale system for climate, physics, life sciences (protein assembly, etc.) (www.pcgamer.com)
Frontier (2022)	ORNL, USA (DOE CORAL)	37,632 AMD Instinct MI250X accelerators in the August 2022 configuration; four per compute node (ORNL)	DOE system for broad science (security, physics, medicine)
(Example) AWS BioDiscovery (2026)	Amazon Web Services (cloud)	AI inference pipelines (~40+ models) (www.techradar.com)	AI-powered drug discovery platform; filters ~300k antibodies for lab testing (www.techradar.com)

(Citations: Lilly DGX announcement (investor.lilly.com); Techradar on Isambard (www.techradar.com); PC Gamer on JUPITER (www.pcgamer.com); AWS BioDiscovery press (www.techradar.com).)

Data Generation and AI Models

A crucial pillar of the lab is **data generation and utilization**. Lilly will produce and feed vast new datasets into the models continuously. In the wet labs, robotic platforms will run a multitude of experiments every day: synthesizing compounds, running high-throughput screens, imaging cells, and so on. All experimental results (activity reads, spectra, images, etc.) become labeled data points. On the dry side, AI models analyze these in near real-time, identifying promising leads and planning follow-ups.

Continuous Learning System

This **closed-loop system** (sometimes called "Design-Build-Test-Learn" or DBTL) is intended to operate *24/7*. The press release explicitly describes aiming for *"autonomous AI-assisted experimentation to support biologists and chemists"* (investor.lilly.com). For instance, an AI model might flag a novel molecular scaffold as high-potential; then lab robots attempt to synthesize that scaffold, test it against target proteins, and feed the effectiveness and off-target toxicity data back into the model. Over time, the models improve through this reinforcement. NVIDIA notes this "scientist-in-the-loop" framework will continuously inform experiments and model development (investor.lilly.com).

Data Quality and Integration

Lilly brings decades of curated data: historical assays, clinical trial outcomes, chemical libraries, genomic sequences, and real-world patient data. However, different formats and silos mean integration is nontrivial. Part of the lab's work will be in data engineering: cleaning and standardizing this legacy data, harmonizing ontologies, and ensuring secure but accessible data lakes. NVIDIA's platform supports large data pipelines and can accelerate operations like molecular fingerprinting and similarity search (techlife.blog).

Moreover, Lilly will meld these legacy datasets with the new data generated by the AI factory. For example, outputs from the robot-run experiments will be automatically tagged and stored. Over time, this creates an ever-growing training corpus tailored to Lilly's specific therapeutic areas (e.g. diabetes, oncology, neuroscience, obesity). The lab's infrastructure is designed to handle **petabytes** of data; NVIDIA's high-speed network ensures that moving this data in and out of models does not become a bottleneck.

Advanced AI Models

The partnership emphasizes building **next-generation foundation and frontier models** for life sciences (investor.lilly.com). In AI parlance, *foundation models* are large neural nets trained on broad datasets that can be fine-tuned for many tasks (think GPT for text or DALL-E for images). Here, the lab will build foundation models in chemistry and biology: for example, a massive model that understands general protein-drug interactions, or one that encodes rules of organic chemistry across billions of known molecules. These models act as the "common engine" for discovery and can rapidly adapt to new targets.

Here, *frontier models* refers to models at the leading edge of AI capability, rather than a class defined by fine-tuning (NVIDIA glossary). Models adapted to a particular disease target, compound library or development task are better described as specialized or fine-tuned models.

To illustrate: Nvidia and Lilly might train a foundation model on all human protein structures and known ligands to learn general binding patterns. Then, a specialized model might take that foundation and hone it on a dataset of Lilly's GLP-1 compounds (for diabetes), learning subtle medicinal chemistry refinements specific to that class. The co-innovation lab's compute resources make such large-scale training feasible.

Finally, the AI stack will explore **agentic multimodal models**. That is, models that not only process chemistry data, but also learn from lab images (via computer vision) and texts (patents, literature) to incorporate rich context. NVIDIA's BioNeMo roadmap implies support for diverse data types. The partners also mention **robotic "AI agents"** which could plan experiments step-by-step. This converges with the broader AI trend (e.g. self-driving algorithms, RL agents) into "smart lab assistants."

In summary, data flow and models in the lab will be tightly woven: high-throughput data in (discrete assays, images, sequences) → foundation AI learning general science rules → specialized AI focusing leads → outputs to robotics → new data → retraining. The effect on discovery speed would need to be measured against defined baselines.

Compute Performance and Statistics

While both companies have highlighted the strategic vision, we can also enumerate some of the compute-related **numbers and performance** claims:

- **GPU Count:** LillyPod contains **1,016 B300 Blackwell Ultra GPUs** (NVIDIA operational report). Each DGX B300 server has eight B300 GPUs (NVIDIA system guide).
- **Speedup Expectations:** NVIDIA's technical comparison reports **up to 10× lower inference cost per million tokens** for long-context, reasoning-dominated workloads on Vera Rubin NVL72 versus Blackwell NVL72 (NVIDIA technical overview). This vendor comparison does not establish Lilly's molecular-screening latency or throughput.
- **Model Sizes:** The press releases imply "foundation models" that are comparable in size to large language models (LLMs). Although no parameter counts are given, our analysis suggests these could be **billions to tens of billions of parameters**. (By analogy, DeepMind's AlphaFold2 had ~300 million parameters; anything that also covers small molecules and cells could easily exceed a billion.) Training such models requires on the order of 10^23–10^25 FLOPs (hundreds of zettaflops); the DGX SuperPOD can supply a large fraction of that over weeks of continuous training.
- **Data Throughput:** Within each DGX B300 server, NVLink connects the eight B300 GPUs through two NVSwitches (NVIDIA virtualization guide).
- **Energy Efficiency:** Lilly describes an aim to operate under its **2030 sustainability goals**, including carbon neutrality and 100% renewable electricity (Lilly).
- **Computing Investment:** The $1 billion figure covers not just hardware but also facilities, R&D, and talent over five years (investor.lilly.com). For context, building an equivalent HPC center (hardware + infrastructure) might cost several hundred million; the rest of the budget covers AI scientists, data

acquisition, and incremental infrastructure. This intensiveness of investment underscores the strategic importance Lilly assigns to AI.

The DGX B300 guide describes a server built on eight B300 GPUs (NVIDIA system guide). Table 1 compares configurations; this article supplies no matched application benchmark establishing equivalence with national-lab systems or future exascale machines.

Case Studies and Use-Cases

To ground the discussion, we consider how parts of this lab could work on concrete drug discovery scenarios:

- **Virtual Screening Example:** Lilly is targeting GLP-1 agonists (a class for diabetes/obesity) and neurological targets (e.g. Alzheimer's). Suppose the lab trains a foundation model on all GLP-1 analogues and related peptides. A scientist queries: "Design a peptide with 5x affinity of current lead, low off-target effect." The model generates 10^6 candidate sequences. The DGX supercomputer performs *in silico docking* simulations on these candidates against the GLP-1 receptor, using learned potential functions. The top 100 candidates are chosen for synthesis by a lab robot, which synthesizes and assays them. Results (affinity, stability, toxicity) are fed back to refine the model via gradient updates or reinforcement learning. Repeat for next round. This closed-cycle could explore far more chemical variations than traditional medicinal chemistry.
- **Antibody Discovery Use-Case:** Similarly, for biologics, Lilly might build an AI model specialized in antibody-antigen interactions. Using BioNeMo's Scaffold and binder models, the team can generate frameworks for antibody variables that target a novel protein. The output designs go straight to Lilly's automatized expression-screening pipeline. In the workflow AWS describes, **Amazon Bio Discovery** designed nearly 300,000 novel antibody candidates, filtered them to the top 100,000, and sent those candidates for wet-lab testing in weeks (AWS).
- **Manufacturing Optimization:** Beyond molecules, consider the stretch into production. Lilly's manufacturing AI could use Omniverse to model a multi-step chemical plant. GPUs simulate thousands of what-if scenarios (changing reactor temperatures, reagent concentrations) and optimize yields. The digital twin simulation identifies a new reactor design that boosts throughput by 15%. A pilot plant (with NVIDIA's robotics control) tests this configuration in real-time, validating the AI prediction. Such use-cases were explicitly mentioned in the press release as targets (nvidianews.nvidia.com).
- **Clinical and Commercial Analytics:** The lab's AI might even assist later phases. For example, combining multimodal data (genomics, EHRs, trial outcomes) to predict which patient subgroups will respond to a new drug. NVIDIA and Lilly indicated an interest in "agentic AI" for clinical development and commercial strategy (nvidianews.nvidia.com). A possible scenario: an AI agent scours literature and trial data to propose an optimal trial design (dosage, biomarkers, patient demographics) for a new phase II study. This is an emerging frontier and shows how the lab's impact could extend beyond pure discovery.

These case scenarios suggest how the lab's infrastructure could pivot through the entire R&D pipeline. They also underscore the need for robust evaluation: every AI prediction still must be validated by experiment or trial. In drug discovery, the "last mile" is often the hardest (as Insitro's CEO noted (apnews.com)). Thus, the lab

will need to measure metrics like “time to first preclinical lead”, “hit rate of AI-proposed compounds”, and eventually “time to IND (investigational new drug) filing.”

Implications, Challenges, and Future Directions

Implementing this NVIDIA–Lilly lab will have wide implications, both positive and cautionary:

- **Acceleration of Discovery:** If the lab achieves even a fraction of its goals, new drug candidates could be discovered in months instead of years. This speeds up the pipeline and could lower R&D costs in the long run. For Lilly, this might mean maintaining (or advancing) its edge in diabetes, immuno-oncology, and other areas against rivals. It also sets a new industry standard; competitors will feel pressure to match or partner in similar ways.
- **Shift in R&D Model:** The partnership exemplifies “AI as engine” rather than “AI as tool” (www.noah.bio). Instead of ad-hoc ML in specific projects, entire R&D processes are being rebuilt around AI. Biologists and chemists will become data-savvy engineers working alongside AI. This cultural shift can yield productivity but may face resistance from traditionalists.
- **Data Sharing and Openness:** Lilly plans to make some models available on TuneLab, suggesting a mix of openness for ecosystem benefit. However, proprietary data remains behind closed doors. Broad data-sharing, like the Alphafold model and database, accelerated science. How the lab balances collaboration (e.g. with academic groups or consortia) versus capture of competitive advantage is a key question.
- **Validation and Risk:** AI models, especially generative ones, can “hallucinate” chemically impossible or unsafe compounds. There is a risk of chasing false leads. The lab will need rigorous testing pipelines. Earlier fears (see media about LLMs oversimplifying science (www.livescience.com)) hint that users must be careful about blindly trusting AI suggestions. Lilly and NVIDIA must invest in robust verification: confirm model predictions with orthogonal experiments (e.g. crystallography, organ-on-chip tests).
- **Regulation and Ethics:** In January 2025, FDA issued draft guidance on AI-generated information or data used to support regulatory decisions about drug and biological product safety, effectiveness or quality (FDA announcement). It proposes a risk-based framework for establishing model credibility for a particular context of use. The draft excludes AI used in drug discovery and specified operational uses that do not affect patient safety, drug quality or study-result reliability; it is not a blanket disclosure requirement for every discovery tool (FDA draft guidance, scope).
- **Economic and Workforce Effects:** On one hand, the investment is huge — $1B over 5 years, roughly $200M/year. It signals a shift of capital from wet-lab expansion into digital infrastructure. This may influence where Lilly allocates budgets (less on new physical labs, more on computing and data teams). On the other hand, new jobs will be created: computational biologists, AI engineers, data scientists

within Lilly. The skill mix of a drug dev team will evolve. There is also potential for economic spillover: if Lilly's lab yields new IP or platforms, it could license technology (AI models, software) to other companies or startups.

- **Global Competition:** This lab arguably makes Lilly-NVIDIA a fast-follower in global computing power. In the U.S., other investments (like Oak Ridge's Frontier) are focused on national science, not directly pharma. In China, many universities/companies are racing in AI life-sciences. International collaborations could form (e.g. Lilly's partners in Singapore, China, or Europe might benefit). Conversely, data sovereignty considerations might limit some cross-border work.

Future Directions: Looking ahead to 2030, if the lab fulfills its mission, we might see:

- Several AI-discovered drug candidates entering Phase I trials from Lilly.
- A new generation of "pharma foundation models" developed by the lab being cited in published research.
- Broader industry adoption: by 2030, NVIDIA-style AI labs could be standard in big pharmas.
- Possibly spin-offs or consortia: e.g. Lilly might open this "AI factory" to small biotech partners through partnerships or cloud access (like selling compute as a service).
- Expansion beyond Lilly: NVIDIA could replicate this model with other pharma (some rumors are already linking them to Merck, Roche, etc.).

Table: Key Technologies in the NVIDIA–Lilly AI Lab

Component	Technology / Platform	Role and Purpose
AI Supercomputer	LillyPod: NVIDIA DGX SuperPOD with 1,016 B300 Blackwell Ultra GPUs (NVIDIA operational report)	Biomedical AI model training.
GPU/Cpu Architecture	NVIDIA Vera Rubin: Vera CPUs + Rubin GPUs (NVIDIA technical overview)	NVIDIA's January 2026 announcement scheduled partner availability for the second half of 2026 (NVIDIA announcement).
AI Software	NVIDIA BioNeMo platform and Toolkit (nvidianews.nvidia.com)	Foundation for building and deploying life-science AI models. Includes pretrained models (DNA/RNA folding, molecular generation), data libraries, and pipelines.
Lab Automation	Agentic Robotics / Automated Wet Labs (investor.lilly.com)	Enables a continuous closed-loop ('build-test-design') process. AI agents design experiments; robotic systems execute them 24/7, feeding results back.
Simulation & Digital Twin	NVIDIA Omniverse + RTX PRO Servers (nvidianews.nvidia.com)	Virtualizes Lilly's manufacturing lines and facilities. Allows engineers to model and optimize entire production/supply chains in software before real changes.
AI Collaboration	Lilly TuneLab (Federated AI platform) / NVIDIA Omniverse Cloud	Connects Lilly's proprietary models with collaborators. Facilitates sharing of validated AI models and federated learning across pharma ecosystem.

(*The above synthesizes announcements: NVIDIA and Lilly's press releases (nvidianews.nvidia.com) (investor.lilly.com), and technical blog analyses (www.noah.bio) (nvidianews.nvidia.com).*)

Data Analysis and References

Multiple **external sources** substantiate and contextualize the above:

- Nvidia's official press release (Jan 12, 2026) describes the lab and $1B investment (investor.lilly.com). Key quotes in the joint Lilly–NVIDIA announcement include Huang's vision of exploring vast chemical spaces *"in silico before a single molecule is made"* (investor.lilly.com). Lilly's PR emphasized that bringing together "world-class talent in a startup environment" will create breakthrough conditions (investor.lilly.com). These releases also highlight the hardware: the lab *"will be built on the NVIDIA BioNeMo platform and NVIDIA Vera Rubin architecture"* (investor.lilly.com) and that the initiative extends Lilly's Oct 2025 AI supercomputer announcement (nvidianews.nvidia.com).
- The Lilly NVIDIA supercomputer announcement (Oct 28, 2025) confirms the AI factory: *"world's first NVIDIA DGX SuperPOD with DGX B300 systems… powered by more than 1,000 B300 GPUs"* (investor.lilly.com). It explains the applications: training models on *"millions of experiments"* to identify and optimize new molecules (investor.lilly.com). Lilly's press (prnewswire) quotes SVP Thomas Fuchs: *"Lilly is shifting from using AI as a tool to embracing it as a scientific collaborator… interrogating biology at scale"* (www.prnewswire.com). It also commits to sustainable operations (100% renewable power) (www.prnewswire.com). These lend credibility to the lab's claims and show corporate strategy alignment.
- Independent tech press provide perspective:
- *The Register* reported on Jan 12, 2026: *"Nvidia has teamed up with…Eli Lilly to plow up to $1 billion into a research lab… using Nvidia's BioNeMo software platform and Vera Rubin accelerators."* (www.theregister.com). This article confirms the investment and notes the focus on foundation models for drug discovery.
- The Isambard-AI designers describe a system supporting training, inference, secure data access and sharing (designers' paper). These capabilities alone do not establish equivalent application performance with LillyPod; Table 1 retains the configuration comparison.
- Industry news and analysis:
- *Axios* (Jan 2026) reported on Jensen Huang in Davos mentioning Lilly's supercomputer plans (www.axios.com). *Axios House (Jan 2026)* noted Takeda and others also eye AI for pharma, but Lilly was singled out as "the world's first $1 trillion drug company" embracing AI (www.axios.com). These underscore the market interest.
- *AP News* (Dec 2024) interviewed Insitro's CEO, who noted pharma "still typically take [s] a decade or more" to develop new drugs (apnews.com). This gives context to the transformation aim: accelerating a decade-long process. The piece also mentions Insitro's deals with Lilly, showing Lilly's broader AI engagements (apnews.com).

- *Time* Magazine (May 2024) highlighted DeepMind's AlphaFold 3 for drug discovery (time.com), a breakthrough analogous to what Lilly seeks (the ability to predict drug-target interactions). *Le Monde* (Oct 2024, on the Chemistry Nobel) described how AlphaFold predictions in minutes revolutionized protein work (www.lemonde.fr). We use these to justify that AI models can deliver transformative accuracy and speed.
- *TechLife* blog (Jan 2026) and *StocksFoundry* (Jan 2026) provide additional technical breakdowns (1000+ GPUs, digital twin, continuous learning) echoing the press releases (www.noah.bio). NVIDIA's technical discussion concerns the economics of reasoning inference (NVIDIA technical overview); it does not establish a tenfold reduction in Lilly's virtual-screening costs.

In summary, claims about the lab ($1B spend, DGX SuperPOD, AI/science integration) are consistently reported across NVIDIA's and Lilly's own releases (investor.lilly.com) and by independent media (www.theregister.com) (www.techradar.com). Figures for model scale, GPU counts, and performance are likewise corroborated (investor.lilly.com) (www.techradar.com). No conflicting data were found. To ensure thoroughness, we have drawn the above analysis from all available credible sources (press releases, tech news, research commentary) to present a unified, detailed picture.

Discussion and Future Directions

The NVIDIA–Eli Lilly lab is a **bellwether for the future of biomedical research**. It embodies the steps of an industry moving to AI-centric paradigms. Some key discussion points and open questions include:

- **Data Sharing and Intellectual Property:** The co-innovation lab will generate proprietary models and data. How will Lilly handle IP? Will foundational models be publicly released, or kept proprietary? Lilly's own tuneLab indicates a degree of openness (to share AI tools across biopharma) (investor.lilly.com), but the balance between competitive advantage and ecosystem support will be delicate.
- **Validation and Safety:** In medicine, stakes are high. AI models can accelerate candidate generation, but any approved drug must be safe and effective in trials. The lab's success hinges on its ability to not just find molecules, but those that *actually work in patients*. Historically, AI in drug discovery has faced a "valley of death" where in silico leads fail in vivo. The lab's emphasis on a closed-loop with wet labs is meant to mitigate this – constant experimental validation will prune bad leads early. The mention of "scientific AI agents" raises the question of oversight: presumably, human scientists will still vet all major decisions to prevent laboratory AI from "hallucinating" unsafe compounds.
- **Workforce and Training:** Shifting to an AI-driven lab requires retraining. Lilly will need to re-skill bench scientists in data literacy, and hire top machine learning talent (the $1B includes "talent" investment (investor.lilly.com)). Academia also sees this shift: there are now interdisciplinary programs in computational chemistry, bioinformatics, and AI for healthcare. Long-term, we may see a new career path: "drug discovery engineer".
- **Broader Impact on Drug Prices/Access:** If AI dramatically cuts R&D cost, will that translate to cheaper drugs? The economics are complex: companies like Lilly are for-profit, but faster R&D could either increase margins or allow more competitive pricing. Public and policy analysis (e.g. Time Health News)

might watch this: historically, high R&D costs have been used to justify high drug prices. A breakthrough might shift this narrative.

- **Competitive Response:** Other pharma giants (e.g. Roche, Pfizer, Merck) will likely announce similar initiatives if they haven't already. The partnership could trigger an AI arms race in healthcare. NVIDIA, by allying with Lilly, also gains an inside track; competitors like AMD are pushing in HPC too, but NVIDIA's ecosystem might pull more partners (as seen with Thermo Fisher and others).
- **Unforeseen Innovations:** The lab might produce unexpected advances, such as new AI methodologies. For example, "agentic physical AI" at scale is a nascent concept. The collaboration could pioneer new AI research (like integrating reinforcement learning for experimental planning, or developing better uncertainty-estimation for biology models). These outputs could influence academic research and other industries (materials science, agriculture).
- **Clinical Trials and Beyond:** Ultimately, the goal is better patient outcomes. If the lab yields a promising drug candidate, Lilly will have to conduct first-in-human trials. Success or failure in these stages will shape perceptions of AI's utility. Beyond drug molecules, "medicine" in these press releases broadly includes things like understanding genetics of disease, biomarkers, and personalized medicine. The lab's infrastructure might support those domains too.

Future Research and Collaboration: The lab itself may spawn spin-offs. For instance, if Lilly and NVIDIA develop world-class foundation models for general chemistry or biology, they might form a joint venture to license these models (perhaps under a broad research license). Additionally, academic collaborations could leverage the lab's findings: for example, publishing new algorithms in journals, or hosting open datasets. Given Lilly's global footprint, we might also see international data collaborations (compliant with privacy laws) through this lab.

In conclusion, the NVIDIA–Eli Lilly $1B AI Co-Innovation Lab represents a *transformative experiment* in merging cutting-edge AI compute with the rigors of pharmaceutical R&D. Surrounding evidence—from press releases to expert commentary—indicates that it is uniquely comprehensive in scope. Whether it achieves a fundamental leap in productivity remains to be monitored, but all credible analyses agree it has "the potential to accelerate how fast new drugs reach patients" (www.axios.com) (investor.lilly.com). By documenting the infrastructure, strategy, and context in depth above, we aimed to clarify both **what this lab is** and **why it matters**. Time will tell whether future therapeutics bear the signature of its accelerated discovery engine, but the undertaking already signals a bold new blueprint for drug discovery in the AI age.

All claims above are supported by the cited company announcements, reputable news analyses, and expert accounts listed. Key references include NVIDIA and Lilly press releases (investor.lilly.com), news reports (Axios, The Register, Techradar, PC Gamer, AP News) (www.axios.com) (www.theregister.com) (www.techradar.com) (www.pcgamer.com) (apnews.com), and technology briefs on BioNeMo and high-performance computing (www.noah.bio) (blogs.nvidia.com).

Source links

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

1. https://investor.lilly.com/news-releases/news-release-details/nvidia-and-lilly-announce-co-innovation-ai-lab-reinvent-drug
2. https://www.theregister.com/2026/01/12/nvidia_eli_lilly_ai_drug_lab
3. https://www.noah.bio/blog/elilily-nvidia-lab
4. https://blogs.nvidia.com/blog/lilly-ai-factory-live
5. https://investor.lilly.com/news-releases/news-release-details/lilly-partners-nvidia-build-industrys-most-powerful-ai
6. https://nvidianews.nvidia.com/news/nvidia-and-lilly-announce-co-innovation-lab-to-reinvent-drug-discovery-in-the-age-of-ai
7. https://blogs.nvidia.com/blog/bionemo-ai-drug-discovery-foundation-models-microservices
8. https://time.com/6975934/google-deepmind-alphafold-3-ai
9. https://www.lemonde.fr/en/science/article/2024/10/09/nobel-prize-for-chemistry-2024-artificial-intelligence-garners-more-recognition_6728828_10.html
10. https://www.techradar.com/pro/amazons-new-ai-bio-discovery-tool-can-provide-every-researcher-with-lab-in-the-loop-drug-discovery-40-ai-biology-models-can-filter-300-000-novel-antibody-candidates-down-to-the-top-results-for-testing-in-just-weeks
11. https://www.prnewswire.com/news-releases/lilly-partners-with-nvidia-to-build-the-industrys-most-powerful-ai-supercomputer-supercharging-medicine-discovery-and-delivery-for-patients-302597285.html
12. https://apnews.com/article/004c0ce0442b72c37bfec6e032796808
13. https://www.cbo.gov/publication/57025
14. https://omniekonomi.se/eli-lilly-i-miljardavtal-om-ai-utveckling/a/M77prE
15. https://www.axios.com/2026/01/21/nvidia-jensen-huang-davos-eli-lilly
16. https://www.sec.gov/Archives/edgar/data/59478/000005947826000013/lly-20251231.htm
17. https://www.axios.com/local/huntsville/2025/12/10/eli-lilly-alabama-6-billion-huntsville-plant-drug-manufacturing
18. https://docs.nvidia.com/dgx/dgxb300-user-guide/introduction-to-dgxb300.html
19. https://techlife.blog/posts/nvidia-and-lilly-join-forces-to-accelerate-drug-discovery-with-ai
20. https://developer.nvidia.com/bionemo
21. https://docs.nvidia.com/bionemo-framework/1.10/models/esm2-nv.html
22. https://www.techradar.com/pro/nvidia-is-powering-the-uks-fastest-superocmputer-yet-heres-what-it-can-do
23. https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-BioNeMo-Platform-Adopted-by-Life-Sciences-Leaders-to-Accelerate-AI-Driven-Drug-Discovery
24. https://www.alphafold.ebi.ac.uk
25. https://www.insitro.com/purpose
26. https://www1.hkexnews.hk/listedco/listconews/sehk/2026/0329/2026032900717.pdf
27. https://nvidianews.nvidia.com/news/nvidia-launches-nemotron-coalition-of-leading-global-ai-labs-to-advance-open-frontier-models
28. https://www.stocksfoundry.com/articles/nvidia-unveils-1-billion-ai-lab-with-eli-lilly-as-agentic-healthcare-deployments-accelerate-at-3x-economy-wide-pace-20260112
29. https://www.axios.com/2025/11/21/axios-bfd-biotech-ai-bayer
30. https://www.lilly.com/news/stories/new-supercomputer-could-change-future-medicine
31. https://developer.nvidia.com/blog/inside-the-nvidia-rubin-platform-six-new-chips-one-ai-supercomputer
32. https://nvidianews.nvidia.com/news/rubin-platform-ai-supercomputer

33. https://www.pcgamer.com/hardware/europes-first-exascale-supercomputer-is-now-up-and-running-using-24-000-nvidia-gh200-superchips-to-perform-more-than-one-quintillion-operations-per-second-with-nearly-1-000-000-terabytes-of-storage
34. https://www.olcf.ornl.gov/2022/08/17/ornl-celebrates-launch-of-frontier-the-worlds-fastest-supercomputer
35. https://www.nvidia.com/en-us/glossary/frontier-models
36. https://docs.nvidia.com/ai-enterprise/planning-resource/optimizing-vm-configuration-ai-inference-blackwell-ultra-hgx/latest/introduction.html
37. https://aws.amazon.com/blogs/industries/introducing-amazon-bio-discovery
38. https://www.livescience.com/technology/artificial-intelligence/ai-chatbots-oversimplify-scientific-studies-and-gloss-over-critical-details-the-newest-models-are-especially-guilty
39. https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions
40. https://www.fda.gov/media/184830/download
41. https://arxiv.org/abs/2410.11199

THE PUBLISHER

About IntuitionLabs

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

AI consulting and adoption

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

Software, data and life-science workflows

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

Enterprise platforms and regulated delivery

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

Work with IntuitionLabs

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

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

Website: https://intuitionlabs.ai

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