

AI Across the Drug-R&D Stack

Where a new layer of defensible value is forming across synthesis, structure-activity, data coordination, and consumer health, and how each is moated and valued

Prepared by Elliot Silk | August 2026

Framing

This memo takes a forward-looking view on a key question in AI for Drug R&D: as AI accelerates each step of pharmaceutical development, where does a new and less-consensus layer of defensible value form? The recurring answer across four domains, molecular synthesis, structure-activity prediction, software-enterprise connective tissue, and consumer health, is that value is forming in the data and infrastructure that sit parallel to pharmaceutical assets. The four conclusions that follow are:

1. In the Make step of the Design, Make, Test drug R&D cycle, the edge for closed-loop discovery goes to whoever invests most heavily in physical AI;
2. In structure-activity work, the next AlphaFold-like leap will be an extension of protein-folding models into AI-native *in silico* SAR and, ultimately, *in silico* mouse and human models that de-risk therapeutics before the clinic;
3. In R&D logistics, the durable layer is the AI-enabled connective tissue software platforms, unifying siloed workflows;
4. In consumer health, value accrues to the diagnostic-grade at-home data layer beneath the aggregator, and to the AI engines that turn that data into therapeutics

For each domain I name the actual moat and how the archetype is valued, and I close with a consolidated bull case, bear case, and risk analysis. Figure 1 places the four domains on the drug-R&D value chain.




Figure 1. AI across the Drug R&D Chain and what this memo addresses.

1. AI for the make step: physical AI and lab automation

PERSONAL PERSPECTIVE AND MARKET CONTEXT

Lab automation is roughly \$7-9B in 2026, growing to about \$16B by 2035, with automated liquid handlers still the single largest equipment segment at about 37% of revenue.¹ The incumbent value pool is liquid handling (Hamilton, Tecan, Opentrons); the un-served adjacency is dexterous solid and bespoke-synthesis automation. Solid dosing, powder handling, and milligram-scale dispensing remain largely manual because they require adaptive, human-like manipulation that fixed-deck handlers cannot do.^{2,3} Physical AI and lab automation are more nuanced than a single autonomous-lab story, and the archetypes matter. The self-driving-lab

literature groups the hardware into three designs: fluidic platforms (liquid handling), fixed-robotic-arm workcells that operate standalone instruments (Chemspeed, Automata), and mobile manipulators, free-roaming robots that automate the scientist rather than the bench.⁴ The mobile-manipulator frontier is where the real bottleneck sits: dexterous handling of solids and bespoke chemistry, demonstrated by the Cooper-lab mobile chemist and now pursued by incumbents such as the ABB-Roche physical-AI collaboration, built increasingly on the NVIDIA Cosmos, Isaac GR00T, and LeRobot stack.^{5,6} Layered on top are business-model archetypes: cloud labs sold as a service (Emerald Cloud Lab, Strateos, Ginkgo's Cloud Lab), programmable synthesis (Chemify's chemputation), and fully closed-loop discovery platforms that own the whole method. Autonomy is a spectrum, and most "autonomous" labs today are level two to three: humans still set the loop. I am most interested in the hardware innovators solving the manipulation problem, because the fluidics and fixed-deck layers are commoditized and the software loop is only as good as the physical actions it can take.

Among the closed-loop platforms, Dash Bio, Arctoris, and Lila mark a clean spectrum (Figure 2): Dash productizes development bioanalysis as a fast CRO replacement (\$30M Series A);⁷ Arctoris runs a robotic data foundry that executes partner-designed experiments;⁸ and Lila is the most ambitious and best-capitalized, owning the whole scientific method as a vertically integrated platform at \$550M raised and roughly \$1.3B with NVIDIA's venture arm in the round.⁹ The deeper question is not who automates, but who uses AI to run the physical lab itself, and how heavily: Dash applies AI to throughput and QC, Arctoris applies ML as a decision layer over partner-set experiments, and Lila puts general-purpose lab robots, foundation-scale models, and sensors in one closed loop that generates its own data, reportedly more than 10 trillion tokens of AI-generated experimental reasoning, and retrains on it.⁹

Today, I would bet on Lila. It is building 235,000-plus square feet of robotic "AI Science Factories" and is backed by NVIDIA's venture arm, In-Q-Tel, and ADIA, the heaviest physical-AI commitment among the three.⁹ Underwriting that bet is threefold: it runs general-purpose lab robots whose unit costs have roughly halved over the last decade rather than bespoke foundry rigs; it pairs the hardware with an AI-




Figure 2. The Build layer: Comparison of Lila, Dash Bio, and Arctoris' models.

and-data flywheel that compounds value per cycle rather than selling foundry time as a service; and it targets cross-domain discovery rather than a single cell-engineering market. Owning hardware is a moat only if utilization, general-purpose reuse, and a data flywheel hold, which is the test any closed-loop bet must pass. Recursion (merged with Exscientia) and Insilico Medicine are the other credible full-stack operators, though their centers of gravity are computational biology and generative chemistry rather than general-purpose robotics,^{10,11} while Chemify is the cleanest make-step specialist, turning xDL code into robotic synthesis.¹²

MOAT, VALUATION, AND THE BULL / BEAR

Moat and valuation. The moat here is a proprietary experimental-data flywheel plus hardware and integration know-how, with composition-of-matter IP layered on for players that carry an internal pipeline. These are valued as a hybrid: an infrastructure or deep-tech multiple on throughput and data, plus risk-adjusted NPV on any pipeline, with a scarcity premium for the few credible full-stack operators. Capital intensity is the swing factor.

Bull. Dexterous solid-handling is genuinely unsolved, the make step is the smallest but most structurally mispriced pool, and the heaviest physical-AI bet, Lila today, compounds proprietary data no one else has.

Bear. Autonomous science can be oversold, and capital costs and utilization have sunk foundries before. Ginkgo Bioworks is the cautionary case, a bespoke bio-foundry that over-built, saw revenue fall roughly 49% year-over-year into 2026, cut more than half its workforce, and executed a reverse split, and low-cost self-driving labs show the loop can be built without nine-figure infrastructure.

Key risk. The bet works only if utilization, general-purpose reuse, and a data flywheel hold; that is the underwriting test for any closed-loop hardware platform, Lila included.

2. AI for SAR and therapeutic platforms

PERSONAL PERSPECTIVE AND MARKET CONTEXT

AI drug discovery is about \$5-6B in 2026 and is forecast to reach roughly \$25B by 2035, with lead optimization the largest sub-segment.¹⁴ The question I keep getting asked is what the next AlphaFold will be. My answer starts with a distinction the field routinely blurs: de-risking chemistry and de-risking biology are different problems, with different data, different maturity, and different prizes.

De-risking chemistry, can we design a potent, selective, synthesizable, drug-like molecule, is comparatively tractable because it is grounded in physics and rich assay data. This is where the AlphaFold-like leap is actually landing: the shift from static structure to affinity and property prediction, or fast in-silico SAR. Structure has commoditized, as Boltz-2, Chai-1, and Protenix now ship with open weights and AlphaFold3's were released for non-commercial use, so the model is no longer the moat; **the durable asset is the proprietary affinity and experimental data that trains it.**¹⁵ Boltz-2 added ligand-affinity prediction because structure alone does not indicate binding strength, and the proprietary frontier is Isomorphic's IsoDDE, reportedly beating AF3 by 2.3x and Boltz-2 by roughly 20x on hard antibody-antigen sets, backed by \$600M and Novartis, Lilly, and J&J.^{16,17} In-silico tox and ADMET is Schrödinger's physics-plus-ML territory. Chemistry de-risking is real and compounding, but it mostly compresses the cost and time of lead optimization.

De-risking biology, is this the right target and will it work in a human, is the harder and far larger prize, because roughly nine in ten clinical failures are driven by biology, efficacy and safety, not by the molecule.¹⁸ This is where virtual cells, in-silico mouse and human models, and virtual patients aim, and

where the data is scarce, high-dimensional, and expensive rather than abundant. Progress is real but early: in cardiotoxicity, in-silico models have already outperformed animal studies at predicting pro-arrhythmic risk,¹⁹ and QuantHealth simulates trials in silico for 12 of the top 20 pharma.²⁰ **The company that de-risks biology in silico before the clinic captures the bulk of the value, because that is where attrition and cost concentrate, but that is a data problem no open model solves, and it sits years behind chemistry.**

MOAT, VALUATION, AND THE BULL / BEAR

Moat and valuation. In chemistry de-risking the moat is proprietary affinity and assay data plus validated accuracy and pharma design-wins, since the model itself is commoditizing; these are valued on software and licensing revenue multiples with milestone and royalty upside, with Schrödinger the public comparable. In biology de-risking the moat is a proprietary perturbation and clinical dataset and, for asset-generating players, composition-of-matter IP; these are valued as platform-plus-pipeline (rNPV) with a heavy data-scarcity premium.

Bull. In-silico de-risking attacks the single largest cost in the industry, the roughly 90% clinical failure rate.

Bear. Biology is not chemistry: models are commoditizing on the chemistry side and nowhere near reliable on the biology side, and “virtual human” claims run well ahead of the evidence.

Key risk. Paying biology-de-risking prices for what is really chemistry de-risking, or backing a model with no proprietary data engine underneath it.

3. AI as R&D connective tissue

PERSONAL PERSPECTIVE AND MARKET CONTEXT

Pharma R&D is fragmented across discovery, development, and regulatory functions, with data trapped in silos and legacy ELN/LIMS systems. The opening is software that becomes the system of record and coordination layer across every workflow. Laboratory informatics is roughly \$3B in 2026, the broader AI-enabled drug-discovery software-and-services market about \$5-6B and growing toward \$25B by 2035, and both sit on top of the roughly \$93B CRO market this layer coordinates.^{14,21,22} This archetype is already crowded. Valkai (Sequoia- and Shine-backed, reportedly more than 20,000 users across top pharma) is building the agentic coordination layer across the pipeline;²³ Potato (a Draper-led seed) backs autonomous science agents that design and run experiments;²⁴ Edison and others push the AI-scientist angle; and Benchling remains the incumbent system of record. Underneath, NVIDIA's BioNeMo is the model-distribution layer.²⁵ Bessemer frames the category as an operating system for biopharma, and demand is real, with Pfizer standing up its own autonomous labs.^{26,27}

So what does a winning company look like, and how does it fend off the labs? Two threats define the answer. First, agent quality is converging, so a thin copilot on a frontier model is not a business. Second, the frontier labs are now forward-deploying: OpenAI's Deployment Company and Anthropic's Ode, a \$1.5B joint venture with Blackstone, put engineers inside enterprises to build custom AI into existing workflows, the Palantir forward-deployed-engineer model, and Ginkgo has already run an autonomous-experiment collaboration with OpenAI.^{28,4} A generalist lab with a forward-deployed team can rebuild a shallow workflow tool. **The winner is therefore not the best agent but the system that becomes the regulated system of record and owns the proprietary experiment-and-outcome data that accumulates inside it, carries GxP and 21 CFR Part 11 validation depth a generalist engagement will not match, and earns enough workflow lock-in and network effect across the pipeline that ripping it out is untenable.** The defense against forward deployment is ownership of the proprietary data and process: pharma is already reluctant to hand its proprietary workflows to OpenAI or Anthropic precisely because doing so

lets those labs compete with it, which is why many are keeping that capability in-house and why a trusted, neutral, compliant vertical system of record still has room. Forward-deployed services are project-based and do not own the data layer; the durable company does.

MOAT, VALUATION, AND THE BULL / BEAR

Moat and valuation. The moat is user-stickiness and switching costs as the system of record, compounding into a proprietary workflow-data and network-effect moat, with compliance depth as a secondary barrier. These are valued as enterprise SaaS on ARR multiples with net revenue retention and seat or usage expansion, at a premium for data-network effects and a discount for anything that is a thin model wrapper.

Bull. Whoever becomes the R&D system of record owns a compounding data asset on top of a roughly \$93B CRO base.

Bear. The field is crowded, agents are commoditizing, and the frontier labs can forward-deploy against shallow tools.

Key risk. Backing a copilot with no data-substrate or compliance moat, which is exactly what forward deployment erases.

4. AI in consumer health

PERSONAL PERSPECTIVE AND MARKET CONTEXT

At the patient-facing end of the stack, the AI-native aggregator captures value while point services and a fast-growing stream of personal data become the substrate beneath it. US telehealth is roughly \$44-65B in 2026 and the global market about \$88B, growing double digits toward the low hundreds of billions by the mid-2030s.²⁹ Hims is the clearest aggregator: it turned partners like Prenuvo and Eight Sleep into discount benefits (\$400 off a scan) and expanded to ten partners including Flo,^{30,31} and raised \$870M partly for an AI-native diagnosis-treatment-delivery stack, the tell being that the platform commoditizes each point service into a feature.³² **The more interesting shift is the data itself: roughly 40% of Americans now wear a Whoop, Apple Watch, Garmin, or Oura, generating continuous biometric data that mostly sits outside HIPAA.**³³ That data is monetized unevenly, since Apple and Garmin lean on hardware and subscriptions and do not sell health data, while Oura and Whoop share aggregated, de-identified data with research and pharma partners, and both face 2026 privacy litigation over biometric handling.³⁴ In practice the data is collected at scale but under-integrated and only partially monetized. The forward-looking opportunity is twofold: a normalization layer that unifies fragmented wearable data, and the shift from correlation-grade signals to diagnostic-grade at-home measurement, where a wave of newcos sits, including Rythm Health (needle-free at-home blood testing, \$79 a month, 25-plus biomarkers)³⁵ and Trace Health (pre-seed, urine biochemistry), a category with comparables such as Hormona and Eli Health.³⁶

The strategic question is what AI does with this data once it is diagnostic-grade and longitudinal. The most instructive example sits one layer deeper: NewLimit runs a proprietary AI model (Ambrosia) over custom aging assays and single-cell profiling to screen transcription-factor combinations, testing more than 3,000 to find some 20 sets that restore youthful function in aged cells, then retraining on the data each experiment generates.³⁷ It raised a \$435M Series C at a \$3.1B valuation on preclinical data alone, precisely because the compounding asset is the proprietary biological data feeding the model, not any single molecule.³⁸ The open question is whether the consumer and at-home data layer above can become the front end of that same loop: if continuous, individual-level biochemistry were fused with molecular and epigenetic readouts into a full-stack, longitudinal dataset, could an AI model design personalized

therapeutics the way Ambrosia designs reprogramming payloads? Who assembles that dataset, who owns it, and whether value accrues to the data, the model, or the delivery, is the pattern this memo argues across every domain.

MOAT, VALUATION, AND THE BULL / BEAR

Moat and valuation. For aggregators and at-home testing the moat is brand and customer-acquisition efficiency, regulatory clearance (CLIA/LDT), and full-stack integration, with a data-network effect as longitudinal data accumulates; these are valued as consumer subscriptions on LTV/CAC and revenue multiples. For the data-to-therapeutics engines the moat is a proprietary biological dataset plus composition-of-matter IP, valued as platform-plus-pipeline (rNPV) with a data-scarcity premium, which is how NewLimit raised at \$3.1B on preclinical data.

Bull. Continuous, individual-level biochemistry is being generated at scale and is barely monetized; the normalization layer and the diagnostic-grade at-home layer are open.

Bear. Wearable data is noisy and correlation-grade, privacy and regulatory exposure is rising, and consumer acquisition costs are unforgiving.

Key risk. Mistaking a wellness app for an infrastructure or data-asset company.

Moats and valuation by archetype

The table consolidates the actual moat and the valuation basis for each archetype discussed above. The throughline is that the defensible asset is rarely the model or the molecule; it is the proprietary, compounding data underneath, or the composition-of-matter IP where an asset is being generated.

Archetype	Primary moat	How it is valued	Key risk
Physical AI / self-driving labs	Proprietary experimental-data flywheel + hardware/integration know-how (+ pipeline IP)	Infrastructure multiple on throughput/data + pipeline rNPV; scarcity premium	Capital intensity, utilization, unproven autonomy
In-silico chemistry (SAR, ADMET)	Proprietary affinity/assay data + validated accuracy + pharma design-wins	Software / licensing revenue multiples + milestone and royalty upside	Model commoditization as weights open
In-silico biology (virtual cell/patient)	Scarce proprietary perturbation and clinical data	Platform + pipeline (rNPV) with a data-scarcity premium	Biology is early, data-poor, unproven
R&D connective tissue	System-of-record stickiness + proprietary workflow data + compliance depth	Enterprise SaaS on ARR multiples + net revenue retention	Crowding + forward deployment from the labs
Consumer aggregator / at-home data	Brand and CAC efficiency + regulatory clearance + data-network effect	Consumer subscription LTV/CAC + revenue multiples	Noisy data, privacy and regulatory exposure
Data-to-therapeutics engine	Proprietary biological dataset + composition-of-matter IP	Platform + pipeline rNPV with a data premium	Preclinical-to-human translation risk

Bull case, bear case, and key risks

Bull case. Across all four domains, AI is opening a new, ownable layer in the data and infrastructure each industrialized step produces, sitting alongside the asset rather than replacing it. The make step is the smallest but most structurally mispriced pool; chemistry de-risking is a compounding software moat; the R&D system of record is a durable data-network business on a roughly \$93B CRO base; and the diagnostic-grade at-home data layer is unowned. Assets still create billion-dollar value, but the next durable layer is the data engine underneath each step, and in small-molecule chemistry that engine is barely built.

Bear case. The consensus “autonomous everything” narrative is priced ahead of the evidence. The most-hyped full-stack labs are capital-intensive and unproven on utilization; models are commoditizing on the chemistry side and unreliable on the biology side; the connective-tissue field is crowded and exposed to the frontier labs' own forward deployment; and consumer data is noisy with rising privacy and regulatory risk. Several of these categories will compress to low-margin services or be absorbed by incumbents and the labs.

Key risks and how I would underwrite them. First, capital intensity and utilization in physical AI: prefer the hardware-manipulation layer and make-step specialists over the most-hyped science factory, and underwrite on utilization, not narrative. Second, data moat versus model: back the proprietary data engine, not the model, since weights are opening. Third, chemistry versus biology: do not pay biology-de-risking prices for chemistry de-risking. Fourth, forward deployment: require a system-of-record data moat and compliance depth, not a thin copilot. Fifth, regulatory and privacy exposure in consumer data. The throughline is one underwriting test in every domain: does the company own a compounding, proprietary data asset that a better model, a cheaper robot, or a forward-deployed lab team cannot replicate?

References

- (1) MarketsandMarkets; Straits Research. Lab Automation Market, 2026 (~\$7-9B in 2026; ~\$16B by 2035; liquid handlers ~37% of revenue).
- (2) Lunt, A. M.; et al. Modular, Multi-Robot Integration of Laboratories: An Autonomous Workflow for Solid-State Chemistry. *Chem. Sci.* 2024, 15 (7), 2456.
- (3) SCU-Hand: A Funnel-Shaped Robotic Hand for Milligram-Scale Powder Handling. *arXiv:2512.07091*, 2025.
- (4) Autonomous Chemistry: Navigating Self-Driving Labs (three hardware archetypes: fluidic, fixed-arm, mobile; autonomy levels; low-cost SDLs; Ginkgo Cloud Lab + OpenAI collaboration). *ScienceDirect*, 2024; QPillars, Self-Driving Labs in 2026, 2026.
- (5) Self-Driving Labs Are Changing How Chemists Work (Cooper-lab mobile chemist). *C&EN*, 2026; ABB-Roche Physical AI in Laboratory Settings. *MarketScale*, 2026.
- (6) NVIDIA. Physical-AI Stack: Cosmos World Models, Isaac GR00T, LeRobot. *nvidianews.nvidia.com*, 2026.
- (7) BioSpace. Dash Bio Raises \$30M Series A to Accelerate Drug Development Timelines, Aug 2026.
- (8) BioPharmaTrend. Case Study: Arctoris Ulysses, End-to-End Wet-Lab Automation with ML, 2022.
- (9) Lila Sciences. AI Science Factories: General-Purpose Lab Robots + Foundation-Scale Models (\$550M; ~\$1.3B; ~10T tokens of AI-generated reasoning; 235,000+ sq ft; NVentures, In-Q-Tel, ADIA). *Excedr*; *Sacra*; *Contrary Research*; *Metatrends*, 2025-2026.
- (10) Recursion. Recursion-Exscientia Merger; NVIDIA-Powered BioHive-2 (up to 2.2M experiments/week). *ir.recursion.com*; *GEN*, 2024.
- (11) Sacra. Insilico Medicine: Generative Chemistry, Robotics-Integrated Lab, and Internal Clinical Pipeline, 2026.

- (12) Chemify. Chempuration: χ DL Code to Robotic Synthesis; ~\$107M raised (Wing Venture Capital, Insight Partners). C&EN; BusinessWire, 2025-2026.
- (13) Ginkgo Bioworks. 2026 Financial Results and Restructuring (revenue down ~49% YoY; workforce reduced >50%; reverse split; refocus on autonomous-lab services). Investing.com; StockTitan, 2026.
- (14) Fortune Business Insights; Roots Analysis. Artificial Intelligence in Drug Discovery Market, 2026 (~\$5-6B in 2026; ~\$25B by 2035).
- (15) Comparison of AlphaFold 3, Boltz-1/Boltz-2, Chai-1, and Protenix (open weights; Boltz-2 ligand affinity). Medium, Mar 2026; Pharma Insight Lab, note.com, May 2026.
- (16) Isomorphic Labs. The Isomorphic Labs Drug Design Engine Unlocks a New Frontier Beyond AlphaFold (IsoDDE). isomorphiclabs.com, Feb 2026.
- (17) Isomorphic Labs Drug Design Engine (\$600M; Novartis, Lilly, J&J). ResearchGate summary, Mar 2026.
- (18) Wong, C. H.; Siah, K. W.; Lo, A. W. Estimation of Clinical Trial Success Rates and Related Parameters. Biostatistics 2019, 20 (2), 273 (~90% clinical attrition, largely efficacy and safety).
- (19) npj Digital Medicine. AI-Driven Virtual Cell Models in Preclinical Research (in-silico cardiotoxicity outperforming animal models), Dec 2025.
- (20) HIT Consultant. QuantHealth Secures \$45M for AI Clinical Simulation (in-silico trials; 12 of top 20 pharma), Aug 2026.
- (21) Healthcare Foresights. Laboratory Information System Market, 2026 (~\$3.2B in 2026).
- (22) MarketsandMarkets. Contract Research Organization Services Market, 2026 (~\$93B in 2026).
- (23) Valkai. Agent Systems for Life Sciences (20,000+ users across top pharma; Sequoia Capital, Shine Capital, Atria Ventures). valkai.com; PitchBook, 2026.
- (24) Potato. Autonomous Science Agent Closes \$4.5M Seed Round (Draper Associates). potato.ai, 2026.
- (25) NVIDIA. Generative AI Is Opening the Next Era of Drug Discovery and Design (BioNeMo; Recursion first partner). blogs.nvidia.com, Nov 2024.
- (26) Bessemer Venture Partners. An Operating System for Biopharma (Potato, Convoke). bvp.com/atlas, May 2026.
- (27) Pharmaceutical Technology. Researchers Hand More R&D to Autonomous Labs (Pfizer autonomous labs), Mar 2026.
- (28) TechCrunch. Anthropic's Ode (\$1.5B JV with Blackstone) and OpenAI's Deployment Company: Forward-Deployed Engineering for the Enterprise, May-Jul 2026.
- (29) Grand View Research; IBISWorld. Telehealth Market, 2026 (US ~\$44-65B; global ~\$88B in 2026).
- (30) Hims & Hers. Introducing Hims & Hers Benefits (Prenuvo, \$400 off). investors.hims.com, Mar 2026.
- (31) Hims & Hers. One Subscription, Ten Ways to Feel Your Best (ten partners incl. Flo). news.hims.com, May 2026.
- (32) Zacks / Yahoo Finance. Hims & Hers Advances AI-Powered Personalized Healthcare (\$870M convertibles; AI-native stack), 2025.
- (33) Electronic Frontier Foundation. Most Smart Watches, Rings, and Bands Lack Basic Transparency and Privacy Features (~40% of Americans own a wearable; data outside HIPAA). eff.org, Jul 2026.
- (34) Wearable data-sharing analyses (Whoop and Oura share aggregated / de-identified data with research and pharma partners; 2026 CIPA and BIPA litigation). GhostVault; Lucas8; SanoLabs, 2026.
- (35) Rythm Health. At-Home Blood Testing (\$79/month, 25+ biomarkers). rythmhealth.com; f4.fund (accessed Aug 2026).
- (36) Coyote Ventures. Hormone Intelligence: The New Frontier of Biomarker Tracking, 2025; Hormona at-home urine hormone testing (category context for Trace Health).
- (37) NewLimit. Ambrosia AI Model, Custom Aging Assays, and Single-Cell Profiling (3,000+ transcription-factor combinations). longevity.technology; Contrary Research, 2025-2026.

(38) NewLimit. \$435M Series C at \$3.1B Valuation (Founders Fund; Eli Lilly Ventures), on Preclinical Data. STAT; Angel Investors Network, Jun 2026.