Vital Signs: AI Drug Discovery

Target discovery, generative chemistry, and pharma partnering

HealthSeed

HEALTHSEED INTELLIGENCE

4 August to 2 September 2026 • Published for the HealthScale community and HealthSeed ecosystem partners

EDITION 001 · MONTH OF SEPTEMBER 2026 · THIS EDITION'S SIGNAL

Pharma spent August buying the tools. The one AI discovery company that reported a profit made it from molecules.

Bristol Myers Squibb signed two AI discovery agreements in sixteen days and disclosed no money for either, while Insilico Medicine reported half-year revenue of US$106.3 million with US$2.7 million of it from software.

This is the premiere issue of Vital Signs: AI Drug Discovery. Every month it covers what is moving in target discovery, generative chemistry, and pharma partnering: the deals pharma is signing, the evidence that the methods work, the money funding them, and the rules governing how any of it reaches a regulator. Europe and Switzerland lead, with the United States and Asia carried where they bear on a European builder.

This track sits further from most readers' daily work than the others, so every specialist term is explained where it first appears. Every claim names an actor, carries a date, and links to a primary source. Where a figure is published by a company about itself, we say so in the line where it appears. Where a finding comes from a preprint that has not been peer reviewed, we say that too.

This edition covers 4 August to 2 September 2026 in four movements: partnering, the benchmark, money, and the rules. Four pharmaceutical companies bought access to somebody else's models this month and published no economics for any of it. In the same weeks, the first blinded prospective test of whether those models beat conventional laboratory work returned an answer with two halves, and a half-year result showed where the revenue in this category is actually coming from. The document governing how that evidence reaches a regulator came into force on 23 July, twelve days before the window opened.

IN THIS EDITION

About this publication

This is the premiere issue of Vital Signs: AI Drug Discovery.

Every month it covers what is moving in target discovery, generative chemistry, and pharma partnering: the deals pharma is signing, the evidence that the methods work, the money funding them, and the rules governing how any of it reaches a regulator. The field is global and the coverage follows it, wherever the work is happening.

This edition covers 4 August to 2 September 2026 in four movements: design, the bench, money, and the rules. Most coverage of this field assumes the models are the hard part, and this month took that apart. Design capability turned out to be purchasable, and improving fast enough that a general-purpose model beat the specialists on their own measure. What none of it could settle was which designs were good, and settling that still takes a laboratory, which is where the cycle's largest European round went and where the regulator's questions now point.

Executive summary

Priority	Signal	Action required
ACT NOW	ICH M15 has been in force since 23 July and names AI and machine learning models in scope	Read it against your own model documentation now, because the alternative is finding the gap during a submission review
ACT NOW	A general-purpose model reached hit rates two to three times the stated industry norm on targets it was never trained for	Reprice what you pay a specialist design vendor for, because the differentiated part may have moved
MONITOR	511 submissions from 29 organisations produced designs that beat the laboratory and no method that won across all three tasks	Name the task your method won and say which one it does not do, before a buyer finds the second answer first
MONITOR	Adaptyv Bio raised US$40 million for wet-lab throughput and counts Chai Discovery, Roche and Novo Nordisk among its customers	Model your own validation capacity as a constraint with a price, not as a line item
MONITOR	Insilico reported US$103.1 million of half-year revenue from discovery and pipeline work against US$2.70 million from software	Check which of those two revenue lines your own plan assumes before your next board meeting

01 Design

Anthropic's general models designed binders at three times the field's hit rate

A minibinder is a small protein built to stick to one chosen target and nothing else. The measure that matters is the hit rate, meaning the share of designs that actually bind when somebody makes them and tests them, and the honest number in this field has been low.

On 18 August Anthropic reported that two of its general-purpose models designed minibinders against fifteen targets with no specialised biology model in the loop. Of 1,320 designs, 354 were confirmed to bind, covering fourteen of the fifteen targets, at hit rates the company puts between 22.6 and 35.1 per cent against a stated industry norm of 10 to 15 per cent. Against one target the rate reached 40 per cent where human participants managed 3.7. Against maltose-binding protein, none of ninety designs bound anything at all.

The same pattern shows in the month's strongest single design. Aureka Biotechnologies produced an antibody reaching 94.7 picomolar, roughly two thousand times stronger than the one it started from. Picomolar is a thousandfold tighter than nanomolar, and tighter binding is what a therapeutic antibody needs. The experimental control it beat took three months of phage display maturation, the standard laboratory method of evolving a binder through successive rounds of selection. The designed version took under a week.

Bristol Myers Squibb bought design models twice in sixteen days

Four pharmaceutical companies bought access to somebody else's models this cycle, and none of the agreements carries a disclosed number.

Announced	Platform	Buyer	What was bought	Terms
5 August	Schrödinger	Bristol Myers Squibb	Bunsen agentic AI co-scientist, computational chemistry stack, RetroSynth planning	None disclosed
20 August	Chai Discovery	Bristol Myers Squibb	Molecular folding and de novo design models across the discovery portfolio	None disclosed
4 August	Phylo	Chugai Pharmaceutical	Biomni Lab agentic platform, connected to Chugai's own research data	None disclosed
28 July	Phylo	Ono Pharmaceutical	Biomni Lab for Ono's discovery scientists	None disclosed
10 August	Amazon Web Services	Novo Nordisk	Preferred cloud and strategic AI partner, London co-innovation hub	None disclosed
6 August	Evotec	Odyssey Therapeutics	Screening, compound libraries and AI-enabled data science	Milestones only, no amounts

Set that against the older shape of a discovery deal, where a platform designs a molecule and takes contingent milestone payments the industry calls biobucks. Relation Therapeutics and GSK announced one on 30 July worth up to US$110 million. Aqemia and Sanofi have run another since 2023 worth up to US$140 million. Both predate this window, and both carry a number.

Sources: Anthropic, 18 August 2026, a company research report that has not been peer reviewed. Schrödinger, 5 August; Chai Discovery, 20 August; Phylo, 4 August and 28 July; Novo Nordisk, 10 August; Evotec, 6 August; Relation Therapeutics, 30 July; Aqemia, 22 July. Every deal here was announced by the platform side, and three of the six carry no separate confirmation from the pharmaceutical counterparty.

THIS IS THE NUMBER TO HOLD UP IN YOUR OWN BOARD PACK

Twenty-two to thirty-five per cent, from a model that was never trained on your problem, against a field norm of ten to fifteen.

This one is for the buyer rather than the seller. If a general-purpose model reaches those rates untrained, the question for a research lead evaluating a specialist design vendor is what exactly the premium is now buying: the model, the data behind it, the people who know which targets it fails on, or the capacity to test what comes out. Those are four different purchases and only one of them is the model. Ask which one is in the quote before renewal, because the answer has changed since the contract was written.

Applies to you if: you buy or sell molecular design capability, or your plan assumes design is the differentiated part.

02 The bench

Twenty-nine organisations submitted 511 antibodies and a laboratory had to build every one

Almost every claim made for AI in molecular design until now has been retrospective. A model is shown data it has never seen, and its predictions are scored against answers that already exist. A prospective benchmark inverts that: participants submit designs before anyone knows the answer, the organisers build them in a laboratory, and everyone finds out together.

Nature Biotechnology published the first one for antibodies on 19 August. AIntibody, organised by Specifica, an IQVIA business, took 511 submissions from 29 organisations across three tasks: improving the binding strength of antibodies drawn from an early sequencing output, ranking candidates within a structurally related cluster, and designing binding regions for proteins absent from any selection output.

The part that took the time is the part that gives the result its authority. Every submission was made as a full-length antibody and measured under standard conditions by surface plasmon resonance, an optical method for watching two molecules bind in real time, confirmed by a second independent method. Five further assays tested developability, meaning the practical properties that decide whether an antibody can be manufactured and dosed at all rather than merely bind in a tube.

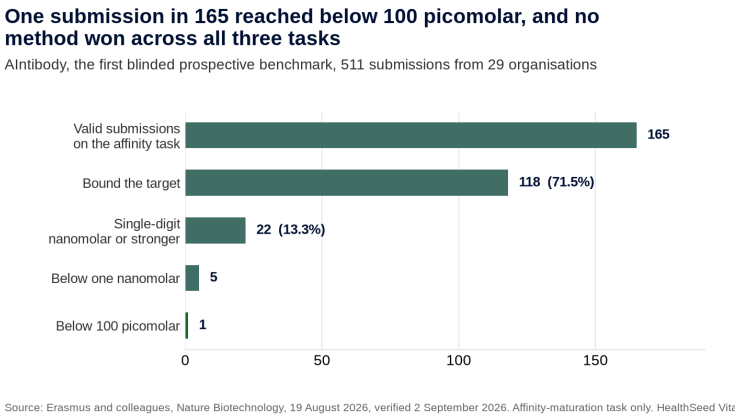

Figure 1. Almost everything bound and almost nothing was strong. The benchmark's authority rests on every one of these having been physically built and measured. Source: Erasmus and colleagues, Nature Biotechnology, 19 August 2026.

Performance varied widely by group and by task, and no single approach dominated. Two limitations belong with those figures. The first round used a single antigen, and it was organised by the same consortium that reported the results, relying on the organisers' integrity for blinding. Neither dissolves the finding, and both are reasons to treat the second round as the one that settles anything.

Aureka beat three months of phage display in under a week, and only the bench could prove it

Aureka took first, second and fifth place. Its winning design reached 94.7 picomolar against a best experimental control of 113 picomolar, and it also produced six antibodies below ten nanomolar that met the developability bar. The figures are published by Aureka and rest on the paper's own measurements.

The comparison is only worth anything because both were measured the same way, in the same laboratory, against the same target. A third benchmark published on 20 August ran twelve molecular generation methods over 176 protein-ligand systems and found architecture-specific trade-offs with no universally best generator. All three point the same way.

Sources: Erasmus and colleagues, Nature Biotechnology, 19 August 2026. Aureka Biotechnologies, 27 August 2026, figures published by the company. University of Texas Health Science Center at Houston, bioRxiv preprint, 20 August 2026, not peer reviewed.

WHAT WE CANNOT YET ASSESS

There is no way to know in advance which method suits your target. That is what "no method won across all three tasks" means, and it has a consequence the coverage has not drawn out.

The experiment is not a validation step that happens after the design. It is the thing that tells you which design to believe. Anyone budgeting for computational design as the expensive part and testing as the follow-on has the cost structure the wrong way round, and the second AIntibody round is the evidence that would change this reading. Until it reports, treat any single-platform performance claim as unranked against its peers.

Applies to you if: you make or evaluate a performance claim for a computational design method.

03 Money

Adaptyv raised US$40 million to test other people's designs in Lausanne

Adaptyv Bio, the laboratory that built and tested Anthropic's designs, announced a US$40 million Series A on 25 August led by Highland Europe with Ace Ventures, ByFounders and Y Combinator. It reports laboratory throughput up more than fivefold in a year and more than a hundred customers, among them Chai Discovery, Boltz, Roche and Novo Nordisk. It is doubling its team in Lausanne and opening a London laboratory in the fourth quarter.

The company states the thesis in its own words: artificial intelligence can only move as fast as the experiments behind it.

Novo Nordisk buys models from Amazon and bench time from Adaptyv

That is the cleanest evidence in the cycle that the two are separable purchases. On 10 August Novo Nordisk named Amazon Web Services its preferred cloud provider and strategic AI partner, opening a London co-innovation hub built on Amazon Bio Discovery and Amazon Bedrock. It is also on Adaptyv's customer list. The same company is buying design capability from one supplier and the capacity to find out whether the designs work from another.

Aureka closed a US$100 million Series B on 10 August, nine days before it won the benchmark, and directed the money at its closed-loop experimental system alongside larger models. Its own account puts the laboratory inside the platform rather than downstream of it.

Insilico earned US$103.1 million from molecules and US$2.70 million from software

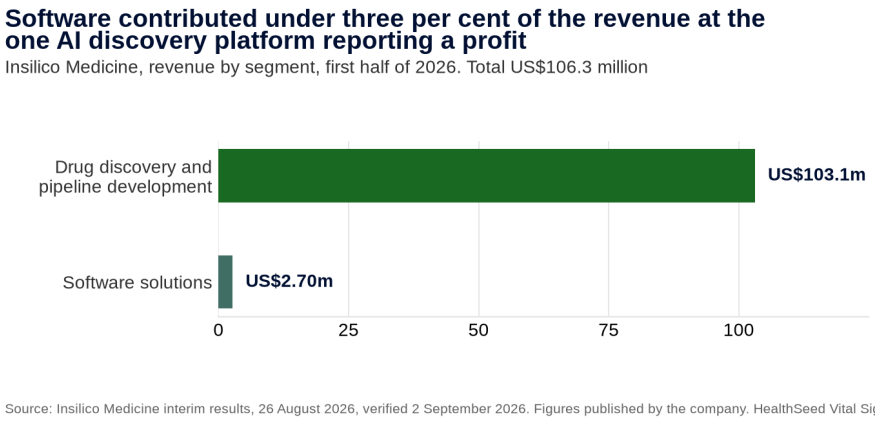

Figure 2. At the one company in the category reporting a profit, the revenue followed the physical work. Source: Insilico Medicine interim results, 26 August 2026, figures published by the company.

Insilico Medicine reported on 26 August that revenue for the first half of 2026 reached US$106.3 million, up 287.2 per cent year on year at a gross margin of 90.3 per cent, with a net profit of US$35.54 million and cash and investments of US$584.8 million. Drug discovery and pipeline development contributed US$103.1 million of that and software solutions contributed US$2.70 million. Recursion, on 5 August, reported quarterly revenue of US$7.67 million against US$19.22 million a year earlier and cut its 2026 cash operating expense guidance to below US$375 million.

We report the European picture as this cycle found it, which is thin. Adaptyv's round is the only substantial European financing in the category inside the window, and Sightera Biosciences, an Antwerp university spin-off, raised EUR 3 million in July, outside it.

Sources: Adaptyv Bio, 25 August 2026; Novo Nordisk, 10 August 2026; Aureka Biotechnologies, 10 August 2026; Insilico Medicine interim results, 26 August 2026; Recursion Pharmaceuticals, 5 August 2026; QBIC, July 2026. Every revenue, cash and round figure is published by the company it concerns. Insilico's half-year and Recursion's quarter are different periods and are not a comparison. A US$73 million Cradle financing reached this cycle's collection dated 24 August 2026; the round was announced on 26 November 2024 and is not reported here.

HEALTHSEED PERSPECTIVE

The case against reading anything into this is the strongest version of the picks-and-shovels argument, and it is usually right. Every technology cycle produces a moment when the tooling looks like the safe bet, and capacity businesses commoditise: throughput becomes cheap, competitors arrive, margins compress, and the interesting economics end up back with whoever owns the asset the capacity was serving. A wet laboratory is a wet laboratory. Grant all of that.

What makes this different is what Adaptyv's customer list actually says. Chai Discovery and Boltz build the models. Roche and Novo Nordisk buy them. Both ends of the market are sending their designs to the same bench, which means the bench sits at the point where a design becomes a fact, and that position is not the same thing as capacity.

The recommendation is for whoever sets your technology strategy, and it is a question rather than an instruction: what do you own, and what are you renting? Renting design capability is now sensible and cheap. Renting the loop that tells you which of your designs was right means renting the only part that generates proprietary data. In our own business development work in this field, the platforms that hold a pharmaceutical partner past the first programme are the ones that arrived with an experimental result the partner could not have produced alone, and that takes nine to eighteen months to build from a standing start.

Applies to you if: you are raising in European AI drug discovery, or you are deciding what to own and what to rent.

04 The rules

ICH M15 asks what your model was validated against, and validation happens at the bench

Every movement above ends in the same place, with a regulator reading a dossier and deciding whether the computational work inside it can be relied on. The document governing that reading came into effect on 23 July 2026, twelve days before this cycle opened, and it belongs in this edition because nothing since supersedes it and because almost nobody in the category has read it.

ICH M15, adopted by the regulatory members of the ICH Assembly on 29 January 2026 and implemented in Europe through EMA's Step 5 publication, sets out general principles for model-informed drug development. Its scope names artificial intelligence and machine learning explicitly, alongside population pharmacokinetics, physiologically based pharmacokinetics, exposure-response analysis, model-based meta-analysis, quantitative systems pharmacology, agent-based models, and disease progression models.

It asks six questions of all of them.

The question	What it wants
Question of interest	The specific decision the model is being used to inform, written down before the model is run
Context of use	Where the model sits in the submission and what else supports the same conclusion
Model influence	How much of the decision rests on the model rather than on experimental data
Consequences of a wrong decision	What happens to a patient if the model is wrong in the direction it is most likely to be wrong
Model risk	Influence and consequence taken together, which is what sets the evidence burden
Model impact	What the model actually changed about the development programme

The guideline names overfitting as the specific concern for AI and machine learning, and requires assessment of robustness against data, parameters, assumptions and uncertainty. It does not cover the technical detail of how a model is built. Every one of those six answers is written in experimental evidence rather than in model architecture.

Three further pieces landed either side of it. EMA researchers surveyed 273 regulators, industry professionals, patients, academics and healthcare workers in June 2026 and ranked ten regulatory science priorities for AI across the medicine lifecycle; accuracy and reliability of AI tools came first with every respondent group, by a substantial margin. FDA's 2026 CDER guidance agenda, released in February and expressly non-binding, adds planned supplemental guidance on software assurance for AI-based systems in drug manufacturing and clinical investigations. And on 27 August the FDA ran a workshop on model-integrated evidence in generic drug development.

Sources: ICH M15 Step 4 final guideline, 29 January 2026, and EMA Step 5 publication EMA/CHMP/ICH/496426/2024 of 9 February 2026, coming into effect 23 July 2026. The EMA Step 5 PDF did not yield readable text on retrieval, and the detail above is taken from the ICH Step 4 text and from regulatory trade coverage of the adopted guideline, which agree on every date. EMA regulatory science priorities, preprint, June 2026. FDA CDER 2026 guidance agenda, February 2026. FDA model-integrated evidence workshop, 27 August 2026.

THE DOCUMENT TO READ BEFORE YOUR COMPETITORS DO

ICH M15 is free and short, and it already binds the three agencies most of this readership will file with. It does not tell you how to build a model, and that is the point. It tells you what a regulator will ask you to have written down about the one you built, which is a different document from the one your computational team currently keeps.

At a platform company the person who should read it this month is whoever will sign the regulatory sections of your first partnership, because the model risk framework in the table above is what a pharmaceutical partner's own regulatory function will use to score you during diligence. Writing that documentation while the experiments are running costs a few weeks. Reconstructing it eighteen months later, for a model that has since been retrained, is the version nobody has ever finished on time.

Applies to you if: your evidence package contains a computational model that will reach a regulator, directly or through a partner.

Critical deadlines, next nine months

Date	Market	Event and action required
23 July 2026, in force	ICH members	ICH M15 governs documentation of model-derived evidence, AI and machine learning included. **Read it against your own model documentation.**
23 to 27 October 2026	Europe	ESMO Congress, Madrid, where Insilico's first-in-human data for ISM6331 and Generate Biomedicines' GB-4362 abstract are both accepted. **Diarise both before setting oncology comparables.**
First half 2028	United States	Initial data from Recursion's ZINNIA trial, on the company's own guidance, and the earliest controlled efficacy readout on the record for any candidate here. **Set the date any AI-attrition claim is tested against.**

The argument

The case against everything above is that one month of capacity financing is not a structural claim. One Series A in Lausanne, one benchmark on one antigen and one half-year result cannot carry a conclusion about where value sits in a field this young, and a reader who has watched three tooling cycles commoditise is right to want more before moving anything.

What survives that is the direction of travel in every piece of evidence the cycle produced. A model with no biology training reached hit rates the specialists publish as their own. Four pharmaceutical companies bought design capability and disclosed nothing, which is what buyers do when a thing is becoming a commodity. The one prospective test of whether any of it works had to be conducted in a laboratory, cost three months of physical work for its comparator alone, and still could not say which method to use next time. And the regulator's own framework asks six questions that are answered with experiments.

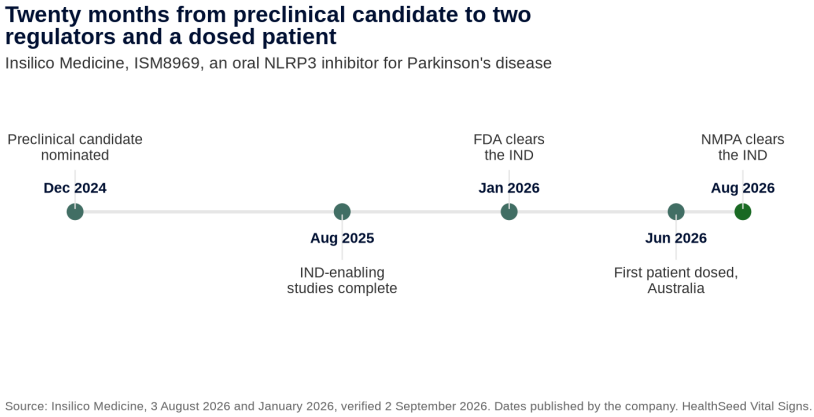

Figure 3. The fastest programme in the cycle, and the design was the short part of it. Source: Insilico Medicine, 3 August 2026 and January 2026, dates published by the company.

The counter-example is worth holding, because it is the fastest thing in the cycle. Insilico took ISM8969 from preclinical candidate to clearance by two regulators and a dosed patient in twenty months, on a molecule its Chemistry42 engine designed. Even there, the design was the short part. Enabling studies, two regulatory reviews and a first-in-human dose took the other eighteen months, and none of that was accelerated by the model that drew the compound.

For a company building here over the next eighteen months, the practical consequence is that the two halves of the work now carry different prices and different risks. Design capability is a purchase, and it is getting cheaper. The loop that tells you which design was right is an asset, and this month it is what raised money, what a benchmark depended on, and what a regulator will ask you to describe.

What do you own, and what are you renting?

Most companies in this category have not written the answer down, and the two lines have separated far enough this month that the answer is no longer obvious. Establishing it is a few days of work against a documented position, and it can be closed before your next board meeting. We are reachable at contact@healthseed.vc.

About HealthSeed

HealthSeed is a Swiss healthcare venture studio and commercialisation partner. We work alongside biotech, medtech, diagnostics, digital health, and AI companies as they enter and scale across Europe, the United States, and the Middle East and North Africa, backed by an Expert Community of more than 35 specialists and a network of partners who have run the functions they advise on. We also build ventures and AI products of our own, so we read this market as participants. Vital Signs is where we publish what we are reading and what we think it means, with a primary source behind every claim.

HealthSeed AG, Weidmannstrasse 5, 8046 Zürich, Switzerland. CH-020.3.056.106-6.

Sources

Every item was opened and verified on 2 September 2026, the send date, rather than on the date it was collected. Four items that reached this cycle's collection were dropped after that verification.

01 Anthropic, autonomous de novo protein design with independent testing by Adaptyv Bio and Twist Bioscience, 18 August 2026, not peer reviewed
https://www.anthropic.com/research/Claude-accelerates-protein-design

02 Schrödinger, strategic collaboration and software agreement with Bristol Myers Squibb, 5 August 2026
https://ir.schrodinger.com/press-releases/news-details/2026/Schrdinger-Announces-Strategic-Collaboration-and-Software-Agreement-with-Bristol-Myers-Squibb-to-Deploy-AI-Co-Scientist-Bunsen-for-Agentic-Drug-Discovery/default.aspx

03 Chai Discovery, collaboration with Bristol Myers Squibb, 19 August dateline, reported 20 August 2026
https://www.businesswire.com/news/home/20260819146905/en/Chai-Discovery-Announces-Collaboration-with-Bristol-Myers-Squibb-to-Advance-AI-Driven-Antibody-Discovery

04 Phylo and Chugai Pharmaceutical, Biomni Lab deployment, 4 August 2026
https://phylo.bio/blog/chugai-partnership

05 Phylo and Ono Pharmaceutical, 28 July 2026
https://www.prnewswire.com/news-releases/ono-pharmaceutical-partners-with-phylo-to-embed-agentic-ai-with-every-discovery-scientist-302835864.html

06 Novo Nordisk and Amazon Web Services, strategic partnership, 10 August 2026
https://www.novonordisk.com/news-and-media/news-and-ir-materials/news-details.html?id=916594

07 Evotec and Odyssey Therapeutics, AI-enabled R&D collaboration, 6 August 2026
https://www.evotec.com/news/evotec-and-odyssey-therapeutics-enter-ai-enabled-r-d-collaboration-in-autoimmune-and-inflammatory-diseases

08 Relation Therapeutics and GSK, expanded collaboration, 30 July 2026, outside the cycle window
https://www.relationrx.com/news/relation-announces-expanded-collaboration-with-gsk

09 Aqemia and Sanofi, target nomination and milestone, 22 July 2026, outside the cycle window
https://aqemia.com/news

10 Erasmus and colleagues, a blinded, prospective benchmark of in silico antibody discovery anchored to experimental affinity and developability, Nature Biotechnology, 19 August 2026
https://www.nature.com/articles/s41587-026-03238-6

11 Aureka Biotechnologies, AIntibody benchmark result, 27 August 2026, figures published by the company
https://www.manilatimes.net/2026/08/27/tmt-newswire/pr-newswire/nature-biotechnology-aureka-wins-the-global-blinded-ai-antibody-benchmark-ai-design-surpasses-the-best-experimental-result/2412994

12 Aureka Biotechnologies, US$100 million Series B, 10 August 2026
https://www.aurekabio.com/news/61

13 University of Texas Health Science Center at Houston, benchmark of twelve molecular generation methods, bioRxiv preprint, 20 August 2026, not peer reviewed
https://www.biorxiv.org/content/10.64898/2026.08.14.744939v1

14 Adaptyv Bio, US$40 million Series A, 25 August 2026
https://www.adaptyvbio.com/blog/series-a

15 Insilico Medicine, interim results for the first half of 2026, 26 August 2026
https://www.prnewswire.com/news-releases/commercialization-roadmap-validated-with-three-digit-million-dollar-revenue-insilico-medicine-announces-interim-results-for-the-first-half-of-2026-302860466.html

16 Insilico Medicine, NMPA IND approval for ISM8969 in Parkinson's disease, 3 August 2026
https://www.eurekalert.org/news-releases/1138579

17 Insilico Medicine, FDA Fast Track designation for ISM6331, 29 July 2026, outside the cycle window
https://insilico.com/news/034jeh2rl1-insilico-medicine-receives-fda-fast-trac

18 Recursion Pharmaceuticals, second quarter 2026 results, REC-7735 IND clearance and the Genentech target option, 5 August 2026
https://ir.recursion.com/news-releases/news-release-details/recursion-reports-second-quarter-financial-results-genentech

19 Generate Biomedicines, second quarter 2026 results and GB-4362 Phase 1 status, 5 August 2026
https://generatebiomedicines.com/news-releases/generate-biomedicines-inc-reports-second-quarter-2026-financial-results-and-provides-business-update/

20 QBIC, pre-seed financing for Sightera Biosciences, July 2026, outside the cycle window
https://www.qbic.be/nieuws/qbic-iii-provides-pre-seed-funding-to-sightera-biosciences

21 ICH M15, general principles for model-informed drug development, Step 4 final guideline 29 January 2026, in effect 23 July 2026
https://database.ich.org/sites/default/files/ICH_M15_Step4_Final_Guideline_2026_0129.pdf

22 EMA Step 5 publication of ICH M15, EMA/CHMP/ICH/496426/2024, 9 February 2026
https://www.ema.europa.eu/en/documents/scientific-guideline/ich-m15-guideline-general-principles-model-informed-drug-development-step-5_en.pdf

23 EMA researchers, regulatory research priorities for AI use in the medicine lifecycle, preprint, June 2026, outside the cycle window
https://zenodo.org/records/20554817

24 FDA, 2026 CDER guidance agenda, February 2026, expressly non-binding
https://www.fda.gov/media/185228/download

25 FDA, workshop on model-integrated evidence in generic drug development, 27 August 2026
https://sbiaevents.com/web20260827/