AI found the drug. Now prove it.

HealthSeed AG took part in a panel discussion on AI drug discovery at the HealthVC Summit in September, alongside a promising selection of AI drug discovery companies: QurieGen, NEBULA, and Ainnocence. This article builds on that discussion to cover what AI has already delivered in drug discovery, where the harder steps still sit, what investors and pharma need to see before capital and partnerships move, where durable value is likely to settle between biotech and pharma, and what each company expects to have delivered by 2031.

Daren Wilson, Co-Founder and Senior Partner, HealthSeed AG · Drawn from the HealthVC Summit panel, University Children's Hospital Zurich, 3 September 2026 · 12 min read

THE ARGUMENT IN ONE DIAGRAM

AI contribution concentrates in the middle of the drug R&D process, and so has the capital

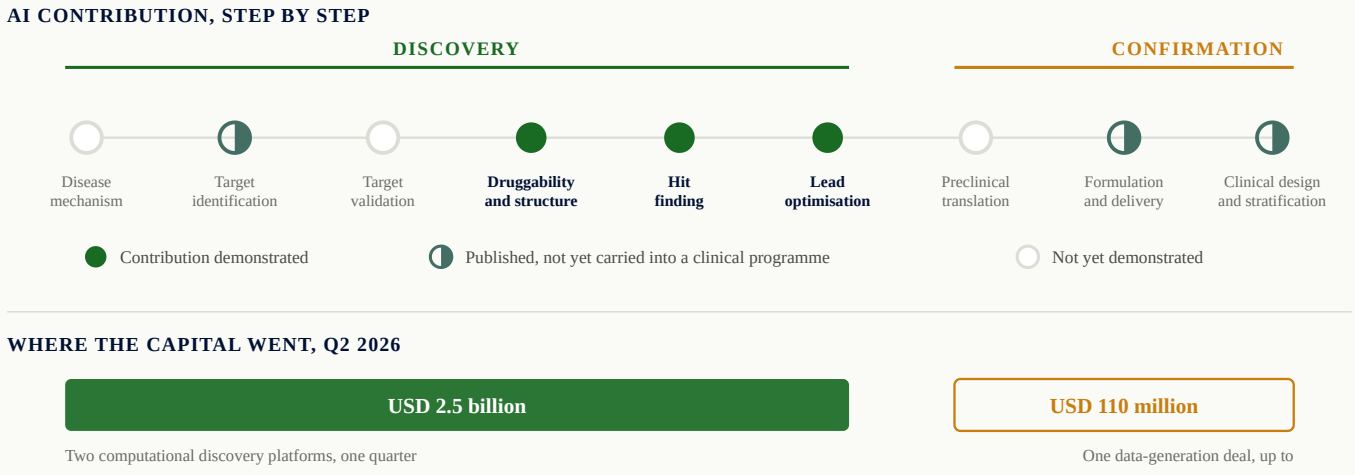

Contribution mapped from Jayatunga et al., Drug Discovery Today, 2024. Capital: Isomorphic Labs and Chai Discovery Q2 2026 rounds; GSK and Relation Therapeutics, July 2026.

I. What AI drug discovery has already delivered

In October 2024 the Nobel committee gave two prizes to the science behind AI drug discovery. The physics prize went to the foundations of neural networks. The chemistry prize went to predicting the three-dimensional shape of a protein from its sequence, a problem biology had not solved in fifty years. Shape matters because a drug works by fitting into a protein the way a key fits a lock, and until 2021 finding that shape meant years of laboratory crystallography for a single protein.

Investment followed the science. Isomorphic Labs raised USD 2.1 billion in a single round in May 2026, and Chai Discovery raised USD 400 million, both to build software that finds and designs drug candidates. More than 170 candidates found with the help of AI are now being tested in patients.

None of them has been approved, and that fact is the one most often used to argue the field has not delivered. It is the wrong test to apply in 2026. A drug that entered its first human trial in 2023 is three years into a process that runs eight to twelve years for every medicine, whichever way the molecule was found. The earliest AI-discovered candidates have not yet reached the point where an approval decision is even possible.

The question that can be answered today is narrower. Which parts of making a medicine has AI made faster, and which parts has it left exactly as they were?

AI has compressed the front end. Choosing which protein to target, and designing a molecule that binds it, now takes months where it took years, and drug discovery pipelines at pharma and biotech companies have roughly doubled in size as a result. Everything after that has stayed where it was. Validation means proving, first in the laboratory, then in animals, then in patients, that the chosen protein really drives the disease, that the drug reaches the tissue where it has to act, and that patients measurably improve. Cell cultures grow at their own speed, animal studies run for months, and clinical trials run for years. No amount of computing power shortens any of it.

Each of the three companies on the panel contributes one or more capabilities to creating a medicine. Thousands of different processes are required to make a medicine, which is why partnering and collaboration are becoming more common in drug discovery. QurieGen measures what a drug does inside a cell, NEBULA finds binding sites on proteins nobody could drug, and Ainnocence designs the molecule itself. We give more detail on each company in Section II.

Five years from now, we will be focusing on how to tailor drug discovery to the patient, and how AI can help find the right combination of treatments for each patient. Getting there means funding validation at the same scale as discovery.

One project shows what the discovery half can now do. Roche and Recursion set out to find which genes drive neurodegeneration. They grew more than ten billion neurons in the laboratory, switched genes off across them one by one, photographed what happened to each cell, and used machine learning to read the results. Nobody decided in advance which genes were worth looking at. A research team testing one hypothesis at a time would need decades to cover the same ground, and would only ever test the hypotheses somebody had thought of.

The result shows up in the pipelines. Beyond the doubling in size, patient populations that were never economically viable to research are now reachable, because the cost of producing a credible starting molecule has fallen far enough to justify programmes that would previously have been dropped. That is a real gain and it belongs to AI drug discovery.

> Up until now, the AI drug discovery space has been well funded for discovery. Take Roche and Recursion. They mapped the whole genome on neurodegeneration off more than ten billion lab-grown neurons, without a hypothesis, and that is something we as humans could never do. Once you get past discovery you have to go to confirmation, and that is the part that still needs the innovation.
>
> Daren Wilson, HealthSeed AG

The clearest evidence that discovery is working came two months before the panel. Insilico Medicine used its own software to pick a target, a protein called TNIK that nobody had run a fibrosis programme against, and a second piece of software to design a molecule that binds it. The result, rentosertib, was tested against placebo in seventy-one patients with idiopathic pulmonary fibrosis, a disease that stiffens the lungs and has no cure. Over twelve weeks, patients on the drug gained 98.4 mL of lung capacity on average while patients on placebo lost 20.3 mL. Blood markers moved in the direction the mechanism predicts. The trial was published in *Nature Medicine* in June 2025 and a Phase III began in July 2026. We looked at what that readout does and does not settle in An AI-designed drug reached Phase III. The evidence bar did not move.

EXHIBIT 1

An AI-nominated target, tested against placebo

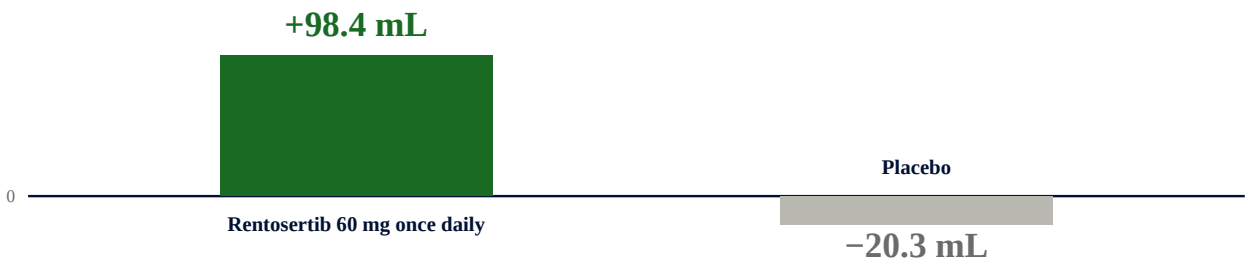

Mean change in forced vital capacity over twelve weeks. 71 patients, 22 sites, randomised, double-blind, placebo-controlled.

Insilico Medicine et al., Nature Medicine, June 2025. Forced vital capacity was a secondary endpoint. Phase III opened July 2026.

Software chose the target, software designed the compound, and patients improved. That is the strongest result AI drug discovery has produced. It is also a single programme in seventy-one patients over twelve weeks, with lung capacity measured as a secondary endpoint and a Phase III that has not reported. Anyone who has run a development programme knows how far that is from an approved medicine. The useful question is what would have to change for the next twenty programmes to look like this one.

II. Three companies, three different missing experiments

Making a medicine runs in a fixed order. First, work out which protein in the body is actually causing the disease. Second, establish whether that protein has a pocket a drug can bind to at all. Third, design a molecule that binds it without also binding a hundred other things. Fourth, show in patients that the result helps.

Three companies presented at the HealthVC Summit, each built against one of the first three steps, and each arguing that a different experiment has not been run. What made the session worth writing up is how precisely their positions fit together.

NEBULA, at the target

Roughly 5,000 proteins are implicated in human disease and around twenty per cent of them can be drugged. The rest carry no visible pocket in the structure a crystallographer solved, so nobody has anywhere to start.

> It sounds like a chemistry problem. It is a physics problem. Those proteins never entered a research programme, and we want to change that.
>
> Mounir Tarek, NEBULA

A protein showing no pocket in a static structure may open one for a few billionths of a second as it moves. NEBULA, which Tarek founded at the end of 2024 out of twenty-five years of simulation work at the CNRS, models that movement atom by atom, detects the openings, and measures how long each one stays open before deciding whether chemistry could use it. Pockets with very short lifetimes get discarded.

EXHIBIT 2

A binding site absent from the crystal structure, recovered from the dynamics

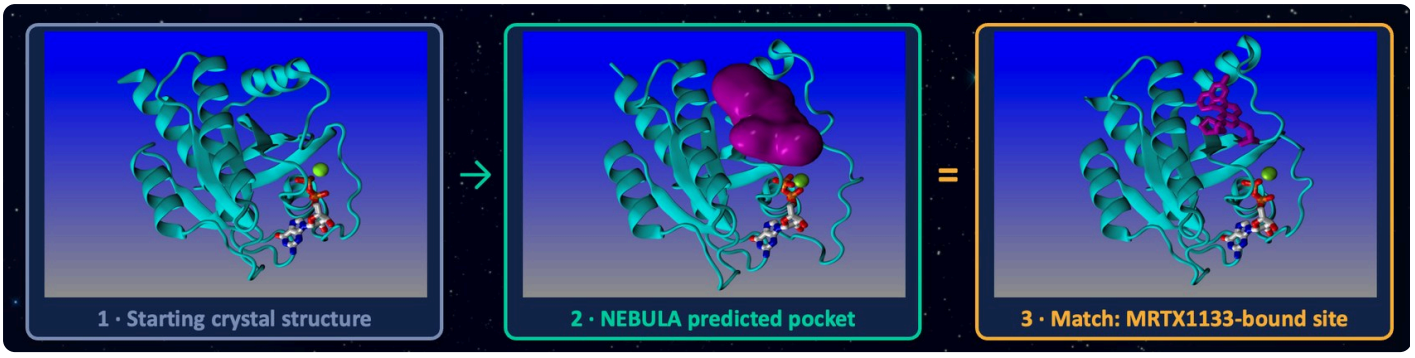

Figure supplied by NEBULA. The algorithm was given one input, the deposited crystal structure, and recovered the site later occupied by MRTX1133. Retrospective.

Their proof point is KRAS, mutated in around twenty-five per cent of cancers and treated as undruggable for forty years. Tarek’s team gave the algorithm one input, the deposited crystal structure of the protein as it was before any inhibitor existed, and nothing else. It recovered the site where marketed drugs now bind.

> It is retrospective. Call it blind rediscovery, weaker than a prediction but stronger than a benchmark score.
>
> Mounir Tarek, NEBULA

He was equally direct about the trade-off, which is the kind of thing a company rarely volunteers in front of investors.

> Unlike the other two companies here, we start again for every target. There is no flywheel. For them, the more data goes in, the better the platform gets. We know there is a price to it. If you are after a large number of assets, we are probably not the ones to provide them.
>
> Mounir Tarek, NEBULA

He also named where his output goes next, unprompted. A novel pocket is an input to a molecule design engine, and he pointed at Ainnocence when he said so.

Ainnocence, at the molecule

Lurong Pan trained as a computational chemist in physics-based methods, hit the same obstacle Tarek describes, and took a different route around it.

> Drug discovery is a multi-target problem, so we needed a method that could screen many targets and see the whole-proteome effect of a therapeutic. Doing that with structure-based methods is almost impossible, so we reduced everything to the sequence level.
>
> Dr Lurong Pan, Ainnocence

The company predicts binders of different modalities against any target from sequence alone, with no structure required. Pan's argument for why this matters is about failure modes rather than speed.

> Most drugs fail because the whole pharmacological profile has to be optimised, and not because of potency.
>
> Dr Lurong Pan, Ainnocence

EXHIBIT 3

Sequence-only design removes the stage that stops many antibody programmes

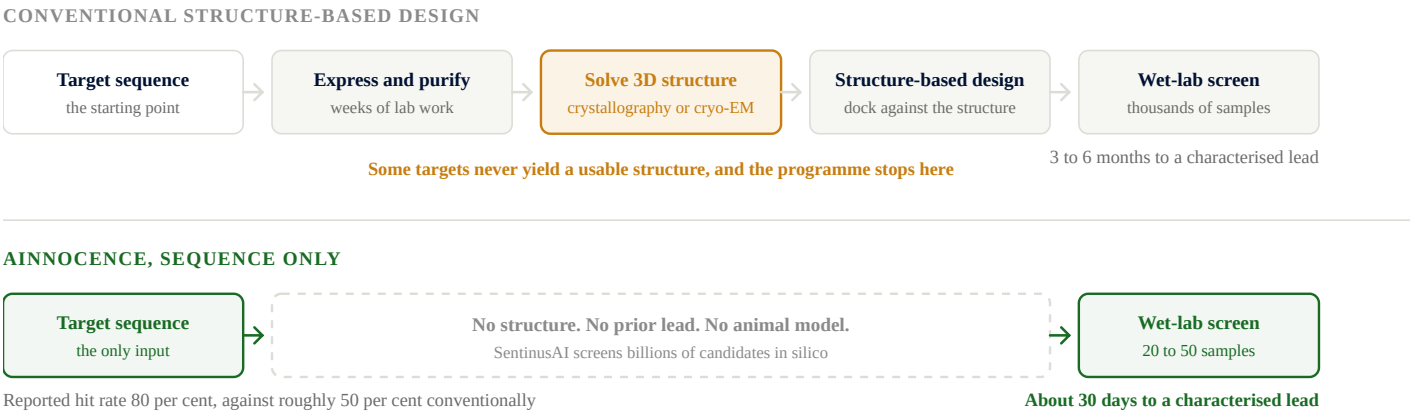

Company-reported figures, Ainnocence. Timeline, hit rate, and sample counts are as published by the company and have not been independently audited.

Ainnocence reports that its SentinusAI platform needs only the amino acid sequence of a target, with no solved structure, no prior lead compound, and no animal model. All figures are company-reported.

Ainnocence runs multi-objective optimisation, scoring every generated molecule across toxicology, ADME, and other physiological outcomes at the same time. The company reports more than a hundred projects in three years, hit rates of ten to sixty per cent in primary screens, around ninety per cent of optimisation projects delivering a preclinical endpoint, and over half of those on targets that were completely novel. It also reports a client molecule now in Phase I. These are company figures and none has been independently audited, which Pan did not dispute.

On the technical claim she was specific. The company's AINN-P1 model carries 167 million parameters and reports the leading stability score on the public ProteinGym benchmark at a Spearman correlation of 0.625, ahead of a 100 billion parameter competitor, and generalises to unseen antibody programmes with an AUC 0.15 higher than a 650 million parameter model roughly four times its size. The work is a preprint and has not yet been through peer review. Publishing the benchmark protocol would let the rest of AI drug discovery check it.

> We are creating the base engine for molecule generation and reducing the cost of wet lab. The next step is enabling the ecosystem that applies it. We cannot do that alone.
>
> Dr Lurong Pan, Ainnocence

QurieGen, at the biology

QurieGen is a Radboud University spin-off from Professor Wilhelm Huck's single-cell lab, and Szollos placed the company deliberately downstream of the other two.

> As scientists, and as humanity, we still understand only a fraction of the biology that drives diseases. That gap is preventing truly transformational innovation in drug discovery and development. I am not talking about generating more selective or more potent molecules faster. I am talking about uncovering novel disease mechanisms and new ways to treat patients, still hidden within the deeper layers of biology.
>
> Jan Szollos, QurieGen

The platform runs perturbation experiments on selected cells and reads the full transcriptome alongside up to 500 membrane and intracellular proteins and phosphoproteins from the same cell, then trains causal models on the result to generate hypotheses about novel targets.

EXHIBIT 4

A live map of the cell rather than a single photograph

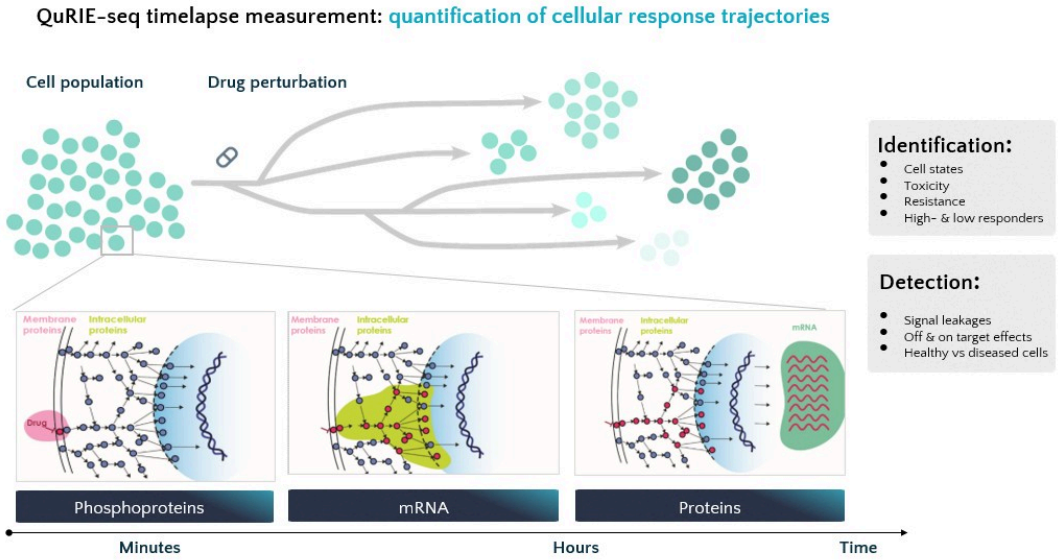

Figure supplied by QurieGen. QuRIE-seq quantifies the full transcriptome alongside up to 500 membrane and intracellular proteins and phosphoproteins from the same single cell, across a time course, so a drug's effect can be followed from intracellular signalling and protein-state changes through to gene expression.

Szollos put the same number on the problem that Tarek did, from the other direction. Around 4,700 proteins are druggable in principle and roughly twenty per cent of that space is being explored. He also said plainly what he has not yet seen.

> AI has helped drug discovery move faster, but I have not yet seen it fundamentally improve our ability to identify truly novel disease mechanisms: new biological ways to treat known diseases that demonstrate real efficacy and meaningfully improve patients' lives.
>
> Jan Szollos, QurieGen

How the three fit together

Their capabilities also run in order. QurieGen's biological measurement tells you which protein is worth attacking. NEBULA's conformational modelling tells you whether that protein can be drugged. Ainnocence's sequence-based design produces the molecule once a pocket exists with no chemistry behind it. Tarek said as much himself on the panel when he named Ainnocence as where his output goes next. Working the other way, spending simulation time on a protein whose role in the disease is unproven is an expensive way to end up with a compound nobody can develop.

None of the three claims to deliver a medicine on its own.

EXHIBIT 5

Three promising AI drug discovery companies, each with a different approach to innovation and technology

QurieGen	NEBULA	Ainnocence
		
Jan Szollos Chief Business Officer	**Mounir Tarek** Chief Executive Officer	**Dr Lurong Pan** Founder and Chief Executive Officer
THEIR APPROACH		
Run biologically relevant perturbation experiments in a systematic way to reveal causal relationships in cells	Use physics to find binding sites on proteins nobody could drug	Design the molecule from sequence alone, no structure required
WHAT THEY HAVE SHOWN		
QuRIE-seq uncovered previously unrecognised cross-talk between BCR and JAK/STAT signalling, revealing JAK1 as an active component of the BCR response	KRAS G12D, binding site recovered from a structure predating any inhibitor	100+ projects, 10–60% primary hit rate, one client molecule reported in Phase I
EVIDENCE LEVEL		
Peer-reviewed publications and 3 industry collaborations	Retrospective	Preprint
IN THEIR OWN WORDS		
"We need to return to the physical world, run the right experiments, generate the right data, and use that data to deconstruct the deeper	"No flywheel. We start again for every target"	"We cannot do that alone"

layers of human biology before training causal virtual cell models.”		
WHAT WOULD STRENGTHEN THE CASE		
Uncover and publish previously unseen causal biological mechanisms that can be explored in a therapeutic context	A preregistered prospective run	Benchmark protocol in the open

Company figures as stated on the panel and confirmed by each company. HealthVC Summit, University Children’s Hospital Zurich, 3 September 2026.

EXHIBIT 6

The fourth step belongs to everyone

Biology	Druggability	Molecule	Validation
What to aim at	Can it be hit	Design the binder	**Prove it in patients**
QurieGen	NEBULA	Ainnocence	**No company on the panel**

Tarek named Ainnocence as the downstream recipient of a novel pocket, unprompted, on the panel.

III. What investors and pharma need to see

The second half of the session turned to evidence, and the bar has moved a long way in five years.

Around four hundred AI drug discovery tools are published free on GitHub. Anyone with the coding skills could assemble a discovery platform from them this year, which means owning software is no longer what makes a company worth backing. Three other things do.

The first is a method nobody else has. Ainnocence works from the amino-acid sequence of a protein and never builds its three-dimensional structure, which is a different way of solving the problem and not a better-tuned version of what everyone else does.

Validated benchmarks come second. Pharma and biotech partners are not there to buy the hype. They are there to buy the evidence, which means published methods, published protocols, and results somebody outside the company can reproduce. Every company on that panel has at least one headline number that currently rests on its own reporting, and each of them would be stronger with the protocol in the open.

Third, and this is the one that has changed most in five years, computational-only is no longer a position. Novelty carried it a few years ago and novelty wears off. Building an AI drug discovery company now means bringing in the wet lab and the validation, and thinking from the bench through preclinical to Phase I and beyond.

Tarek’s response to the same question is the most useful answer a founder can give.

> The only reliable thing is wet-lab data that confirms our prediction. That is exactly what we are raising money for.
>
> Mounir Tarek, NEBULA

> If you are going to a pharma or biotech company, they want to see the evidence. They are not there to buy the hype. They are there to buy the evidence.
>
> Daren Wilson, HealthSeed AG

Szollos named the deal he thinks sets the pattern. In January 2026 GSK licensed Noetik's OCTO-VC virtual cell foundation models for non-small cell lung and colorectal cancer, a five-year agreement worth USD 50 million in upfront and near-term payments plus an annual subscription fee. GSK bought access to a model, not a molecule.

> GSK's licensing of foundation models from Noetik may become a standard model: pharma licenses external AI infrastructure, then fine-tunes it using its most precious resource which is proprietary data from millions of experiments already run internally. This approach becomes even more powerful when you put the model with a lab in the loop.
>
> Jan Szollos, QurieGen

That structure matters because of what it puts a price on. Trained models are becoming easy to obtain. Experimental results from human tissue are not, and a pharma company sitting on decades of its own laboratory data holds something no amount of computing can generate.

IV. Where durable value sits

Will pharma simply build all of this in house? The last two times the industry asked that question it got the same answer.

When high-throughput screening arrived, the assumption was that every large pharma would absorb it and the specialist companies would disappear. They did not disappear. When the genome was mapped, the assumption was that pharma would take genome mapping in house. Many specialists are still out there providing that service, and the same shape is likely here.

Some AI drug discovery companies will sell a service to everyone, which is a durable business and keeps pharma and biotech supplied with capability they do not want to build. Others will hold their method to themselves and develop their own drugs. Both models work, and companies get into trouble when they pitch one and are built for the other.

Pharma will in-house what falls inside its strategic interest, which historically means the data and the clinical pathway, and it is now buying compute on the same logic. Roche has taken on more than 3,500 Nvidia chips to build that capacity.

> The future of drug discovery will be increasingly shaped by biotechs that take risks, explore novel biology, and build specialised platforms. Pharma will increasingly focus on clinical development, regulatory approval, manufacturing, and commercialisation.
>
> Jan Szollos, QurieGen

> When high-throughput screening was developed, you would have thought every pharma company would take it in house. They did not, and many companies still survive doing it. When we mapped the genome, people thought pharma would take over genome mapping. They did not. I think it is the same with AI drug discovery, and business partnering is only going to get bigger.
>
> Daren Wilson, HealthSeed AG

He also made a point about access that deserves more attention than it got in the room. AI is democratising drug discovery. Thirteen years ago pharma held a near-monopoly on it, employing thousands of chemists to fill pipelines. That work moved to biotech, and the tools are now good enough that individual scientists with strong computational skills and enough structural biology to know what they are looking at can do it. Useful molecules will start coming from labs and individuals nobody is currently watching.

The capital implication follows from the same reasoning. Discovery has been well funded. Isomorphic Labs raised USD 2.1 billion in May 2026, around thirty-seven per cent of all digital health funding in the quarter, with Chai Discovery second at USD 400 million. Both are strong companies doing serious science on structure and design.

EXHIBIT 7

Two computational platforms, one data-generation deal

Second quarter 2026. Total digital health equity funding was USD 5.7bn, down 24 per cent quarter on quarter. Sources: CB Insights; company announcements.

WHERE THE CAPITAL WENT, Q2 2026

37%

One round for one computational discovery platform accounted for thirty-seven per cent of digital health funding in the second quarter of 2026.

CB Insights, State of Digital Health Q2 2026. Isomorphic Labs USD 2.1bn Series B, May 2026.

Validation has not been funded at that scale. Wet-lab capacity, getting a drug to the right tissue, and trial design that can detect an effect in the patients most likely to show it are where the returns now sit. The Noetik licence Szollos described, and GSK’s expanded collaboration with Relation Therapeutics in July 2026 worth up to USD 110 million for time-resolved human perturbation data, are early examples of capital finding that layer.

The strongest structural point from the panel is the one every speaker arrived at independently. Nobody can do this alone. Making a medicine is extraordinarily complex, and it now requires several different technologies, several companies, and pharma and biotech working together across each stage. Business partnering is going to get bigger, and the companies that know which part of drug R&D they own will be the ones that partner well. We set out how those deals are actually structured in AI Drug Discovery: What's the Deal?

V. What has to be true by 2031

Each panellist was asked what their company would have delivered in five years, and the answers were unusually concrete.

Szollos wants one number to move. Of the roughly 4,700 druggable proteins, around twenty per cent are being explored. Moving the druggable count from 4,700 to 5,000, and raising the explored fraction by ten per cent, would create medicines for millions of people who have none.

> The low-hanging fruit has been reached. The next frontier is harder: we need to return to the physical world, run the right experiments, generate the right data, and use that data to deconstruct the deeper layers of human biology. The idea that we already have all the data needed to decode every disease, and that the remaining problem is simply more compute and better reasoning, will not get us there.
>
> Jan Szollos, QurieGen

EXHIBIT 8

The number Szollos wants to move by 2031

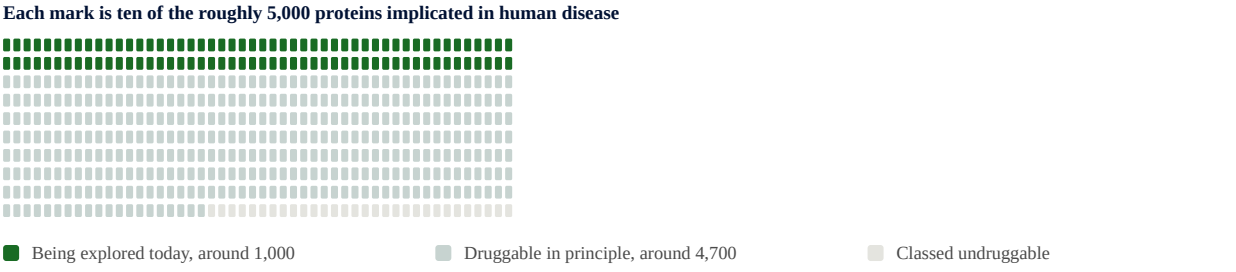

Figures as stated on the panel by Jan Szollos and Mounir Tarek. The 2031 goal Szollos named is moving the druggable count to 5,000 and raising the explored fraction by ten per cent.

Pan framed hers around unmet need. Around fifteen thousand diseases have no standard treatment. Four per cent of the world's population has a rare disease, most of them genetic, most requiring intervention in early childhood. Further pandemics will come and will need a fast response. She believes molecule generation for validated targets is close to a solved problem computationally, and that the constraint is the efficiency of the lab and the clinical trial.

Tarek was blunt about what he needs and what he offers.

> We are not going to do this on our own, because we want to move faster. We can take a few undruggable targets and deliver validated binders within two months. Oncology, central nervous system, inflammation, we are agnostic. Whoever raises their hand first is who we team up with.
>
> Mounir Tarek, NEBULA

That is an open invitation to a pharma partner, made in public, and invitations like it usually go unanswered.

My own answer was about where the attention goes. Five years from now, we will be focusing on how to tailor drug discovery to the patient, and how AI can help find the right combination of treatments for each patient.

Getting there means building out validation, the slower half of drug R&D, with the same seriousness the discovery half received. QurieGen, NEBULA, and Ainnocence each named a piece of what is missing. Whether the capital, the wet-lab capacity, and the partnerships arrive fast enough is what decides how quickly any of it reaches a patient.

The three companies

Each company has reviewed and approved its own entry.

QurieGen

QurieGen is a Dutch techbio company building a “Google Maps for cells” to de-risk and accelerate drug discovery.

We create AI-powered virtual cell models to understand how cells respond to treatment, behave in disease, and transition between states. What sets us apart is our integrated biology-first approach based on proprietary single-cell multi-omics technologies, high-quality in-house datasets, and causal AI models spanning RNA, proteins, phosphoproteins, and epigenetics. Combined with a lab-in-the-loop approach, we map how soft and hard perturbations reshape cellular systems, enabling new target and biomarker discovery and prediction of cellular responses. Kinga Matuła is Chief Executive Officer.

Based Nijmegen, Netherlands **Web** quriegen.com

Open to Pharmaceutical partners for mode-of-action and target identification and validation projects.

NEBULA

Uses physics, with no training data, to find binding sites that no static structure shows, opening targets long written off as undruggable.

A CNRS TechBio spin-off founded at the end of 2024 in Nancy, out of twenty-five years of molecular simulation research. The platform maps a target's full conformational landscape at atomic resolution, identifies cryptic and allosteric sites, and measures how long each opening persists before committing chemistry to it.

Based Nancy, France **Web** nebulabio.tech

Open to A pharma discovery partnership. Validated binders on undruggable targets within two months, therapeutic area agnostic.

Ainnocence

Designs binders from sequence alone, optimising binding, specificity, developability and manufacturability together, with no solved structure required.

Founded in 2021 and headquartered in California. The platform screens billions of small-molecule and antibody candidates at whole-genome scale without 3D structural modelling, and has been applied across more than sixty therapeutic programmes spanning antibodies, small molecules, siRNA, cell therapies, and synthetic biology.

Based San Francisco, California **Web** ainnocence.com

Open to Academic groups, biotechs, and industry partners for joint development, pilot studies, and long-term programmes.

About the author

Daren Wilson, HealthSeed AG

Co-founder of HealthSeed AG, a Swiss venture studio and commercialisation partner working with biotech, medtech, diagnostics, and digital health companies across Europe, the United States, and MENA.

Before founding HealthSeed, twenty years in pharmaceutical commercial leadership at Roche, Novartis, and Merck, including general management roles and global product strategy on major products in oncology and rare disease.

Based Zürich, Switzerland **Web** healthseed.vc

Open to Companies building for drug R&D: AI drug discovery, wet-lab innovation, clinical trial innovation, and the infrastructure around them.

Sources

Sources Panel: “AI Found the Drug. Now Prove It.” HealthVC Summit, University Children’s Hospital Zurich, 3 September 2026. Panellists Mounir Tarek (NEBULA), Dr Lurong Pan (Ainnocence), Jan Szollos (QurieGen), and Daren Wilson (HealthSeed AG). Insilico Medicine et al. A generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis: a randomised phase 2a trial. *Nature Medicine* , 3 June 2025. Phase III initiation announcement, 7 July 2026. CB Insights, State of Digital Health Q2 2026. Isomorphic Labs Series B, 12 May 2026. Chai Discovery Series C, 2026. GSK and Noetik, OCTO-VC foundation model licence, 8 January 2026. Relation Therapeutics and GSK, 30 July 2026. Jayatunga MKP, Ayers M, Bruens L, Jayanth D, Meier C. How successful are AI-discovered drugs in clinical trials? *Drug Discovery Today* 29(6), 104009, 2024. Wang R, Jin K, and Pan L, AINN-P1 protein foundation model, preprint 2026, not peer reviewed. Company figures for NEBULA, Ainnocence, and QurieGen are as stated by each company on the panel and identified as such in the text.