ViPARTNERS

The Next Generation of Drug Discovery

How AI, Physics, and Computational Biology are Reshaping the Discovery of New Medicines

Author: Maria Escala-Garcia, PhD, Investment Associate, Vi Partners · October 2026

Executive summary

Drug discovery has followed the same basic approach for decades: basic research describing a cellular pathway and mechanism of action (e.g. the role of p53 in cell cycle and cancer), identify a target, find a molecule or biologic that engages it, and push the candidate through a long and expensive path to chemistry, manufacturing and controls (CMC) and the clinic. That approach is now under real economic strain. R&D productivity has declined for decades even as the underlying life-science toolkit has improved dramatically, a pattern known as Eroom's Law, Moore's Law spelled backwards: bringing a new drug to market still costs low-to-mid single-digit billions of dollars, takes more than a decade, and succeeds for only a small fraction of candidates that reach human trials. This has started to change at the discovery stage: Insilico Medicine, for example, reports an average of 12 to 18 months from program start to preclinical candidate nomination, against the three to six years that traditional early-stage discovery typically requires.[1] Clinical duration, regulatory review and manufacturing scale-up remain largely unchanged, which is precisely why the early stages are where the economics are shifting.

Computational methods have begun to change the early stages of that economics, and hold the potential to reshape the rest, but not through a single technology. Machine learning, generative AI, physics-based simulation, and increasingly hybrid approaches are being layered across the discovery workflow, from choosing what to work on to predicting how a molecule may behave before it is synthesized. "AI for drug discovery" has frequently become the default label for all of this, but it says little about the different problems these approaches solve, why physics-based methods remain important, or what distinguishes a durable computational discovery platform from an expensive experiment.

The next generation of drug discovery will not be defined by a single winning algorithm. It will emerge from the convergence of data-driven models, generative approaches, physics-based simulation, and increasingly automated experimentation. **The companies most likely to create durable value will be those that combine the right computational method with proprietary data, meaningful experimental feedback, and a business model that captures and compounds the value created.** This paper maps the space around the discovery workflow, then sets out the framework we use to distinguish genuine platforms from bespoke services and from businesses whose technology is interesting, but whose economics are difficult to defend.

USD 2.6bn
Estimated cost per approved drug, including failures (Tufts CSDD)

10–15 yrs
Typical target-to-approval timeline

7.9%
Likelihood of approval from Phase I entry, 2011–2020

12–18 mo
Insilico's average time to preclinical candidate, vs. 3–6 years traditionally

1. Why drug discovery needs to change

In 2012, Jack Scannell and colleagues coined "Eroom's Law," Moore's Law spelled backwards, to describe a persistent decline in pharmaceutical R&D efficiency: the number of new drugs approved per USD 1 billion of inflation-adjusted R&D spend has roughly halved every nine years since 1950, even as sequencing, genomics and high-throughput screening advanced at Moore's-Law-like speed in the same period.[2] Scannell attributed the gap to a "better than the Beatles" problem, in which every new drug is judged against an ever-improving standard of care, and a "cautious regulator" problem, in which the tolerance for safety risk in approved medicines has fallen over time.

Cost estimates vary widely depending on how they are built. A widely cited Tufts Center for the Study of Drug Development estimate puts the figure at USD 2.6 billion per approved drug, including the money spent on candidates that failed and the cost of having capital tied up for a decade. Critics have challenged it because the underlying company data was never published, so the number cannot be independently checked.[3] Deloitte's annual analysis of the top twenty

Eroom's Law: R&D efficiency has halved roughly every nine years

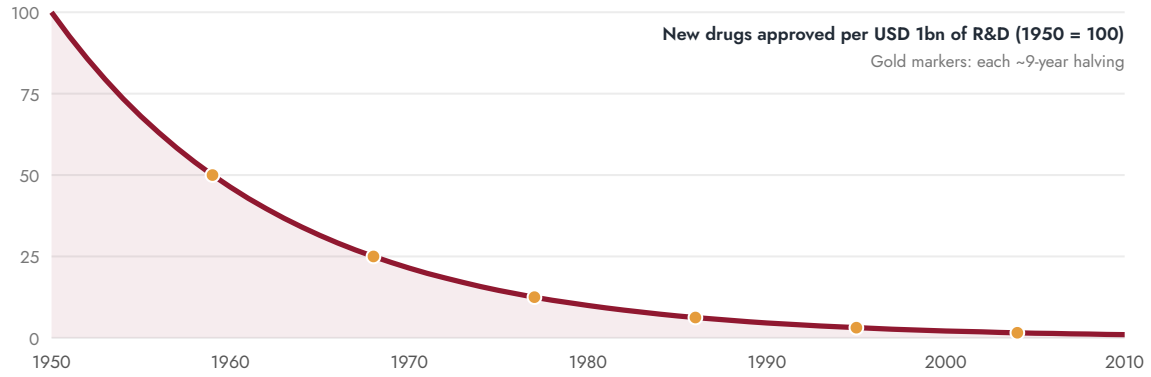

Figure 1 | Illustrative index of new drugs approved per inflation-adjusted USD 1 billion of R&D, drawn from the halving rate reported by Scannell et al.

Source: Scannell et al., Nature Reviews Drug Discovery (2012); Vi Partners illustration.

pharmaceutical companies points the same way: average cost per approved asset was an estimated USD 2.28 billion in 2023 and USD 2.23 billion in 2024, before rising to USD 2.67 billion in 2025.[4] Whatever the precise number, the trend is clear.

Timelines have not improved either. Target-to-approval programs still typically run 10 to 15 years. A landmark analysis of more than 12,700 phase transitions across roughly 9,700 development programs between 2011 and 2020 found the overall likelihood of approval from Phase I entry to be just 7.9 percent, with new molecular entities notably lower at 5.7 percent (small molecules overall at 7.5 percent) and biologics and vaccines somewhat higher, at 9.1 and 9.7 percent respectively.[5] Of the thousands of compounds that enter early screening for a typical program, only a handful become leads, and only a small fraction of those that reach preclinical testing are ultimately approved.

Rising cost, persistently low odds

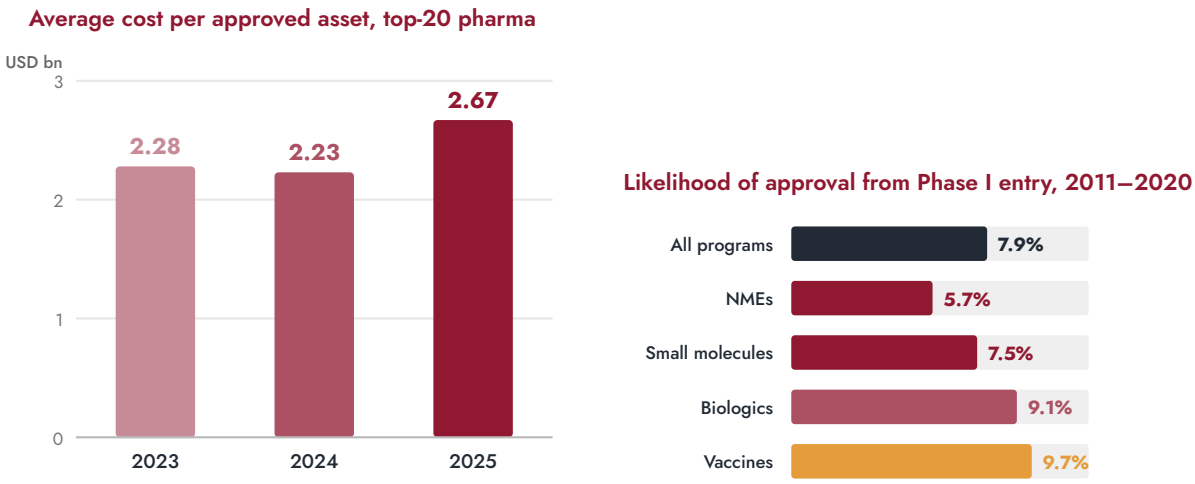

Figure 2 | Cost per approved asset has risen over the period while fewer than one in ten programs entering Phase I reaches approval.

Sources: Deloitte, Measuring the return from pharmaceutical innovation (2023–2025); BIO, Informa Pharma Intelligence & QLS Advisors (2021).

Where the funnel breaks

Programs do not fail evenly across the workflow, and each stage breaks for a different reason. Target identification and validation fail most often not because the chemistry is hard, but because the underlying biological hypothesis does not hold in humans even when it held in preclinical models, a translational failure rather than a technical one. Hit discovery is dominated by false positives and by hits that are real but cannot be developed further, because the molecule resists chemical improvement. Lead optimization is a multi-parameter problem, where improving potency frequently degrades solubility, permeability or safety, and optimizing all of these simultaneously by trial and error is slow. Preclinical in vivo work is meant to catch efficacy and safety risks before human exposure, but its predictive value is inconsistent: reported animal-to-human concordance varies so widely across models and indications that a positive result offers limited advance signal on whether a program will translate.[6]

Clinical translation is where these earlier failures surface at the highest cost: an analysis of Phase II and Phase III failures between 2013 and 2015 found that, among failures with a stated reason, lack of efficacy accounted for 52 percent and safety for 24 percent, with strategic and commercial decisions making up most of the rest, meaning most of what kills late-stage programs is a failure to predict human biology correctly much earlier in the process, not poor late-stage execution.[7] That distinction matters for everything that follows: most of the computational methods in this paper make molecules better (more potent, more selective, safer, easier to synthesize) without necessarily making the underlying target choice, or the trial built around it, any more likely to work in humans. Closing that second gap turns out to be a different problem, which we cover in Section 6.

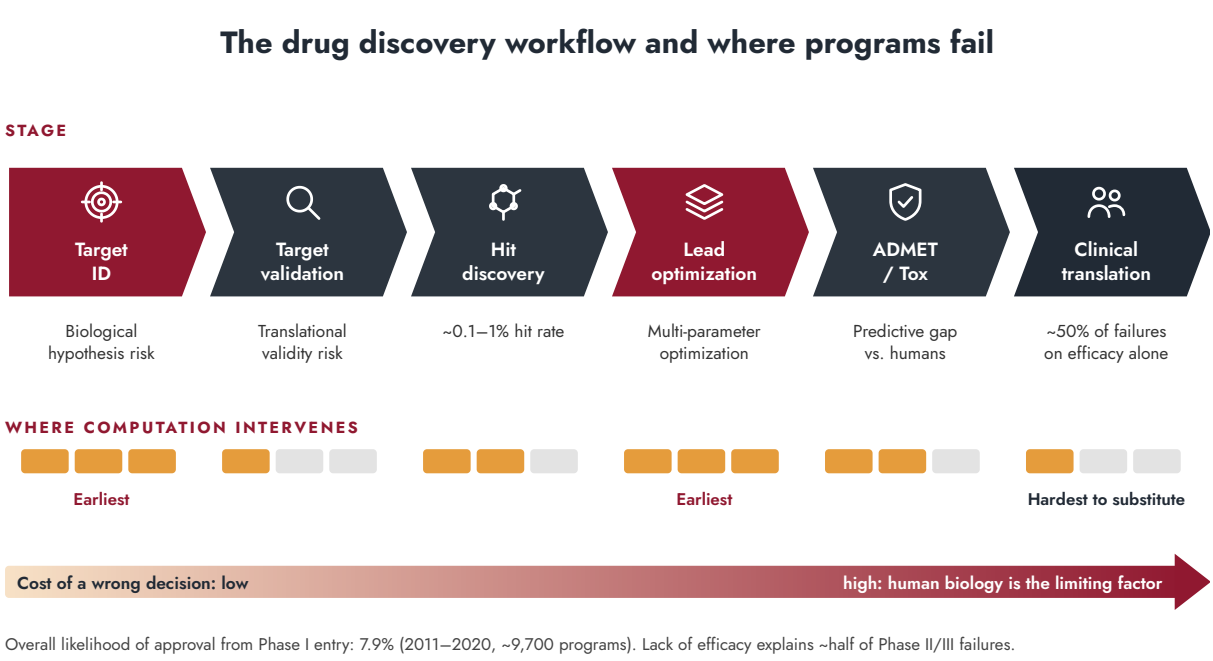

Figure 3 | Computational methods intervene earliest at target identification and lead optimization, where the cost of a wrong decision is lowest, and are hardest to substitute for at clinical translation, where human biology remains the limiting factor.

A related question is worth asking of any platform in this space: which step of the workflow is short on capacity. Target identification and clinical trials are constrained by scarcity: there are relatively few validated targets, and only so many eligible patients who can be recruited, screened and monitored at once, so progress made there tends to compound.

> **Molecule design and optimization face a different constraint: the number of viable candidates a design program can generate has grown faster than the clinical system's capacity to test them, so making that step faster does not by itself move the ceiling on how quickly a therapy can reach approval.**

A platform can be very good at the latter without touching what is actually limiting progress, which is worth asking directly, alongside the science, in any diligence process. We return to this distinction in Section 5.

2. What changed computationally

"AI transforming drug discovery" is the framing everyone reaches for, but it masks an important point: several distinct computational families are maturing in parallel, and their strongest applications increasingly overlap. The relevant question is therefore not which technology wins, but which approach is best suited to the biological problem at hand and how effectively it connects prediction to experiment.

A note on definitions

In this paper, **machine learning** (ML) refers to models that learn patterns from historical data to predict risk, activity, or outcomes. **Generative AI** refers to models that propose new molecules, sequences, or structures rather than only scoring existing ones. **Physics-based methods** compute molecular behavior from underlying physical principles, without requiring training examples of the specific system studied. **Agentic AI** refers to systems that go beyond producing outputs to planning and carrying out multi-step tasks, such as querying databases, running computational tools, or initiating experiments, with limited human input at each step. None of these labels implies the others.

Machine learning
Learns from historical and multi-omic data to prioritize targets, predict activity and estimate risk.

Generative AI
Designs new proteins and small molecules directly, rather than filtering existing libraries.

Physics-based methods
Computes binding and behavior from physical laws; works where training data do not exist.

Hybrid approaches
Physics generates training data and keeps models plausible; learned models make simulation fast enough to scale.

Machine learning from historical and omics data

The oldest and most proven computational family learns patterns from large historical or multi-omic datasets to prioritize targets, predict activity, or estimate risk. Recursion Pharmaceuticals has built its business on billions of cellular microscopy images and transcriptomic profiles, used to infer relationships between genes, compounds and phenotypes at industrial scale. Insilico Medicine's PandaOmics platform identified a novel target for idiopathic pulmonary fibrosis from multi-omic data, and the resulting molecule reached Phase I in under 30 months from program start, a timeline that would be unusual under a conventional discovery process. Their Phase IIa results followed a few years later.[8]

Generative AI: designing rather than screening

A second family designs candidates directly rather than filtering existing libraries. Protein language models such as the ESM family and ProtGPT2 generate novel, natural-like protein sequences from sequence data alone.[9] Diffusion-based tools extend this to joint sequence-and-structure design, generating de novo protein backbones or binders. For small molecules, diffusion and transformer architectures increasingly propose new chemotypes directly, rather than ranking compounds drawn from existing chemical libraries.

Physics-based methods still matter

Physics-based methods remain essential because they do not need training data for the chemotype or target being studied: they compute directly from the underlying physical laws. Free-energy perturbation (FEP) methods can now predict relative binding affinities to within roughly 1 kcal/mol, approaching the reproducibility of the experimental assays they are meant to replace, and underpin commercial platforms such as Schrödinger's FEP+.[10] More detailed quantum-level calculations do better on the hardest targets, such as those built around a metal atom or bound by drugs that form a permanent chemical bond, predicting both how a molecule sits in its target and how tightly it holds. Physics-enabled design has already reached late-stage development: Schrödinger's computational platform contributed to the design of anti-inflammatory drugs like TAK-279 (zasocitinib), a TYK2 inhibitor now under FDA Priority Review in psoriasis following positive Phase III results.

Hybrid approaches: where the field is converging

Neither approach is sufficient on its own: machine-learning models are reliable only on chemistry close to what they have already seen, while physics-based simulation needs no prior data but costs too much to run across every candidate in a large design space. Combining the two, using physics to generate training data and to keep model predictions physically plausible, while using learned models to speed up or fill in between expensive simulations, is where much of the most credible recent work sits. Isomorphic Labs, the DeepMind spin-out built on AlphaFold3, is the clearest large-scale expression of this convergence, pairing deep-learned structure prediction with more classical modeling steps in a single design pipeline. It raised USD 600 million in 2025, led by Thrive Capital, followed by a USD 2.1 billion Series B in May 2026, again led by Thrive Capital, and holds close to USD 3 billion in combined biodollar-heavy partnerships with Eli Lilly and Novartis, alongside a newer, separate collaboration with Johnson & Johnson. In 2025, a free model called Boltz-2, built by MIT and Recursion, came close to matching the accuracy of these physics calculations while running about a thousand times faster. It does not replace physics: it learned from physics and experimental results, and it still needs a predicted structure to work from. But it shows how quickly the line between the two approaches is moving.

2024
Nobel Prize in Chemistry for protein structure prediction and design: the first for an AI-enabled achievement

3M+
Researchers who have used the AlphaFold database

190+
Countries those researchers span

The clearest signal that computational biology has moved from peripheral to central came in October 2024, when the Nobel Prize in Chemistry went to Demis Hassabis and John Jumper for protein structure prediction, the work behind AlphaFold, and to David Baker for computational protein design, the first Nobel Prize awarded for an AI-enabled scientific achievement.[11] The AlphaFold database, built on that work, has since been used by more than three million researchers in over 190 countries.

Where each computational family is most useful

	Machine learning	Foundation models	Physics-based	Generative AI	Automation
Target ID	Core	Emerging	Limited	Emerging	Limited
Design	Supportive	Core	Core	Core	Limited
Optimization	Core	Emerging	Core	Emerging	Limited
Safety / ADMET	Core	Limited	Supportive	Emerging	Supportive
Clinical translation	Emerging	Limited	Limited	Limited	Emerging

Core · Supportive · Emerging · Limited

Figure 4 | A qualitative map across the discovery workflow. No single technology dominates every stage; the strongest platforms combine several, deliberately, rather than defaulting to one.

Source: Vi Partners assessment.

3. Where the platforms are being built

Rather than list companies, we find it more useful to organize the emerging platform landscape into six categories that map onto where in the discovery workflow each one intervenes. The categories are not mutually exclusive, and the strongest platforms increasingly combine several approaches. The examples below are illustrative, chosen to show how different technologies are being commercialized rather than to provide an exhaustive market map.

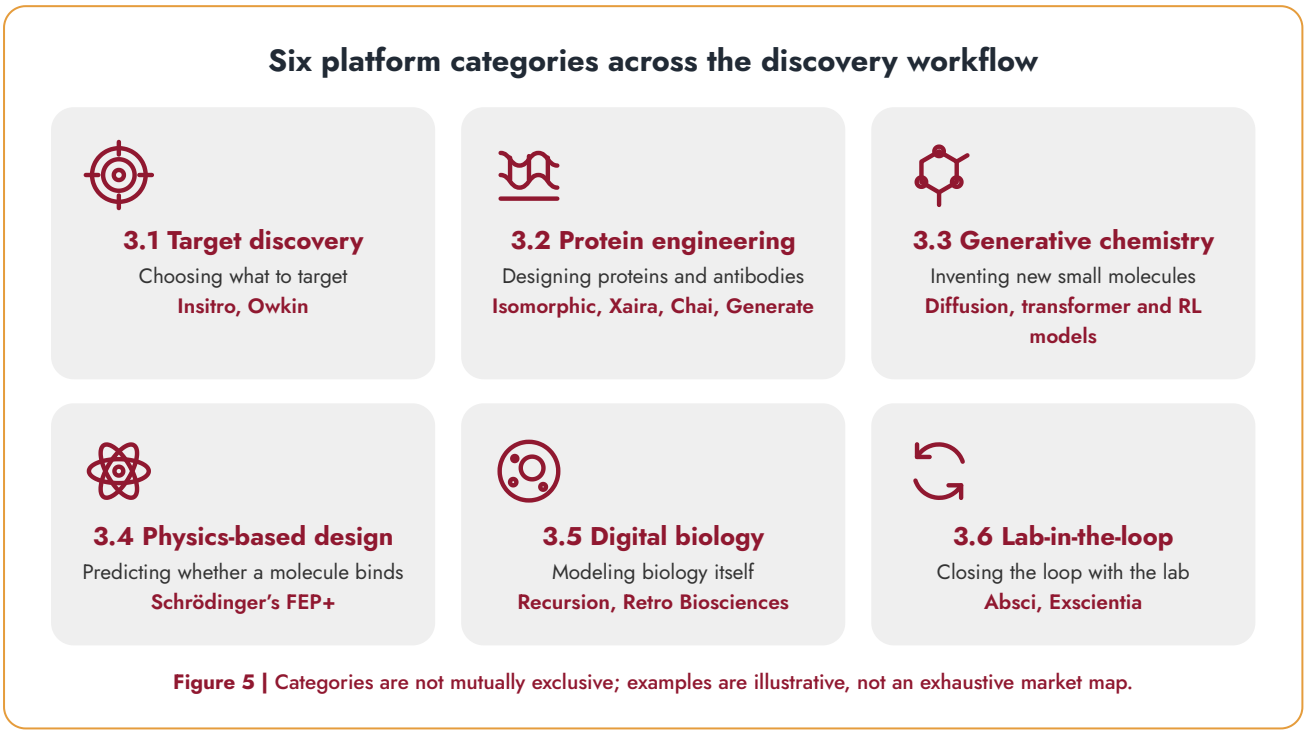

Figure 5 | Categories are not mutually exclusive; examples are illustrative, not an exhaustive market map.

3.1 Choosing what to target (target discovery)

These platforms decide what to go after. Drawing on omics, network biology, causal inference and foundation models trained on multi-omic data, they work out which gene or protein actually drives a disease before any molecule is designed. Insitro raised roughly USD 400 million in a 2021 Series C to combine cellular and clinical data for target identification in metabolic and neurological disease. Owkin, having raised over USD 300 million across its history, licensed its "agentic AI" discovery platform to Sanofi under a five-year agreement in June 2026.

Some target discovery efforts take a different approach to the training population itself: rather than studying an average cohort, they aim to recruit and characterize biological outliers, for example people who resist a disease despite exposure that reliably causes it in others, carriers of a high-risk mutation who never develop the disease, healthy individuals with rare loss-of-function variants, or extreme responders and non-responders to an existing therapy. Comparing outliers to matched, typical-outcome individuals isolates the biology that makes them different. This kind of dataset is inherently rare, whether because it takes years to recruit or because it is already held exclusively by whoever assembled it, making it hard for a competitor to replicate.

3.2 Designing proteins and antibodies (protein engineering)

These platforms design the protein itself, rather than looking for a small molecule to act on one: structure prediction in the AlphaFold3 lineage, protein language models, and de novo protein and antibody design. This is the most heavily capitalized category today. Beyond Isomorphic Labs, Xaira Therapeutics launched in 2024 with more than USD 1 billion in committed capital, led by ARCH Venture Partners and Foresite Capital. Chai Discovery, backed by OpenAI, raised a USD 130 million Series B in December 2025 at a USD 1.3 billion valuation for antibody design from scratch, without target-specific training data or lab iteration. Generate Biomedicines, having raised a USD 273 million Series C in 2023

with Amgen and NVIDIA's venture arm as new investors, completed a Nasdaq listing in early 2026.

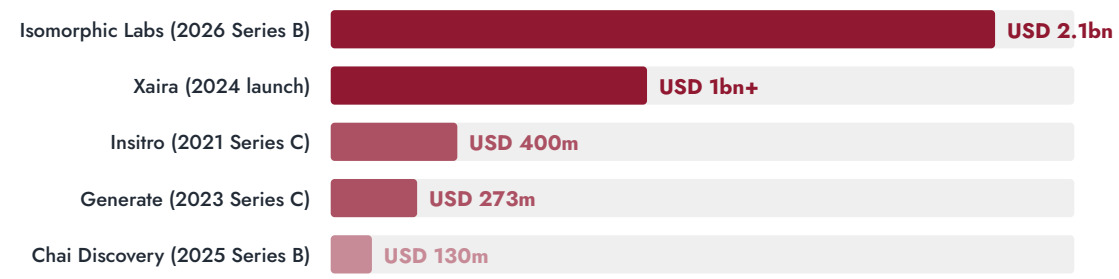

Figure 6 | Protein and antibody design is the most heavily capitalized category today. Rounds as cited in the text; not exhaustive.

3.3 Inventing new small molecules (generative chemistry)

These platforms invent small molecules that have never been made. Generative models built on diffusion, transformer and reinforcement-learning architectures propose the structures, and a second class of model works backwards from a proposal to a practical route for making it from available starting materials. The opportunity is to move beyond screening existing chemical libraries and explore a much larger design space while optimizing several properties simultaneously. The remaining challenge is experimental confirmation: a molecule that scores well computationally still must be synthesized, tested, and shown to behave as predicted. This makes integration with experimental workflows and the ability to learn from proprietary results increasingly important. This category works hand in hand with the physics-led methods described next, which test whether the proposed molecules bind.

3.4 Predicting whether a molecule will bind (physics-based design)

These platforms assess whether a proposed molecule will actually bind its target, and how tightly. Using molecular dynamics, free-energy methods and physics-first binding prediction, they calculate the interaction from physical laws rather than learning it from examples, which is what makes them usable where the data needed to train a model does not exist. Schrödinger's FEP+ remains the commercial reference point. **Sibylla Biotech**, a Vi Partners portfolio company, takes an unusually physics-native approach within this category: rather than working from a protein's final folded structure, it models the short-lived intermediate shapes proteins pass through while folding, using a proprietary method combined with its own AI platform to identify druggable pockets on protein classes that lack one in their final structure. The platform's oncology pipeline includes small-molecule degraders against KRAS and cyclin D1, while a separate collaboration with a pharmaceutical partner applies the same folding-intermediate approach to neuroscience targets.

3.5 Modeling biology itself (digital biology)

These platforms model the biology rather than any single target or molecule, using large-scale imaging and cellular data to build a working representation of how cells behave and respond. Recursion, introduced above, has converted that imaging dataset into commercial validation: more than USD 200 million to date under its Roche/Genentech collaboration and more than USD 130 million under its Sanofi collaboration. Retro Biosciences, backed by Sam Altman, applies a similar large-scale-data approach to cellular reprogramming in pursuit of an explicit goal: extending healthy human lifespan by ten years. It set out to raise USD 1 billion in a Series A in 2025, and in May 2026 announced the initial close of a further round at a USD 1.8 billion pre-money valuation.

3.6 Closing the loop with the lab (lab-in-the-loop discovery)

These platforms close the loop between prediction and experiment: predictions are tested in the lab and the results feed straight back into the model, with as little human delay in between as possible. This design-make-test-learn cycle is

increasingly referred to as lab-in-the-loop discovery. Agentic AI pushes this further: rather than a scientist handing off each step, software agents can select the next experiment, run the relevant models and interpret the results before proposing the following round. Absci has built pharma partnerships around this model worth up to USD 247 million with AstraZeneca, up to USD 610 million with Merck and up to USD 650 million with Almirall, and took a strategic investment from AMD in 2025 to expand its AI-compute capacity. Exscientia ran one of the most visible early versions of this design-make-test-learn loop before combining with Recursion in 2024, discussed further below.

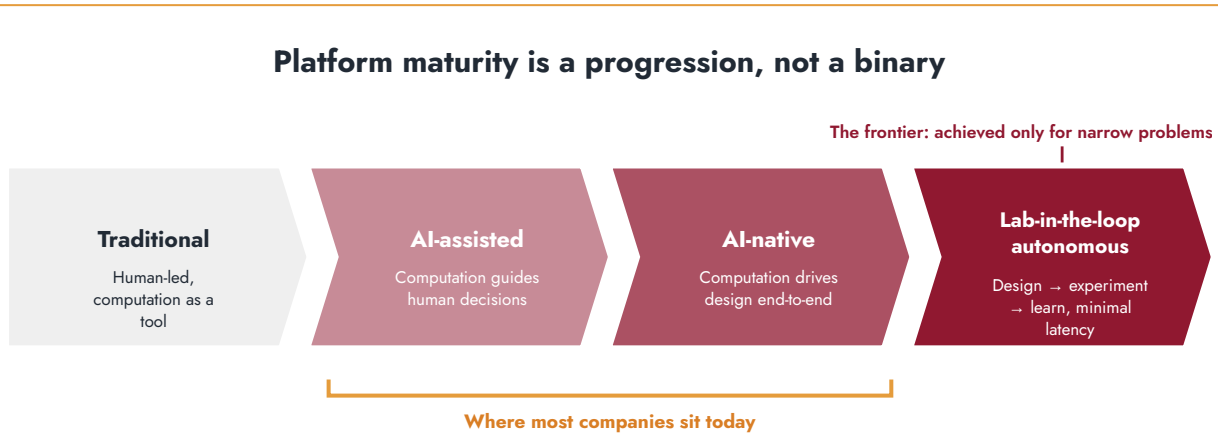

Figure 7 | Most companies today sit between AI-assisted and AI-native. Genuinely autonomous, lab-in-the-loop discovery remains the frontier, achieved so far only for narrowly defined problems.

4. Emerging business models

The platform categories above cut across several distinct commercial models, and the choice of model matters as much as the underlying science for how a company should be valued and de-risked. The key question is not simply who pays, but what the company retains from each program and whether each engagement strengthens the platform over time.

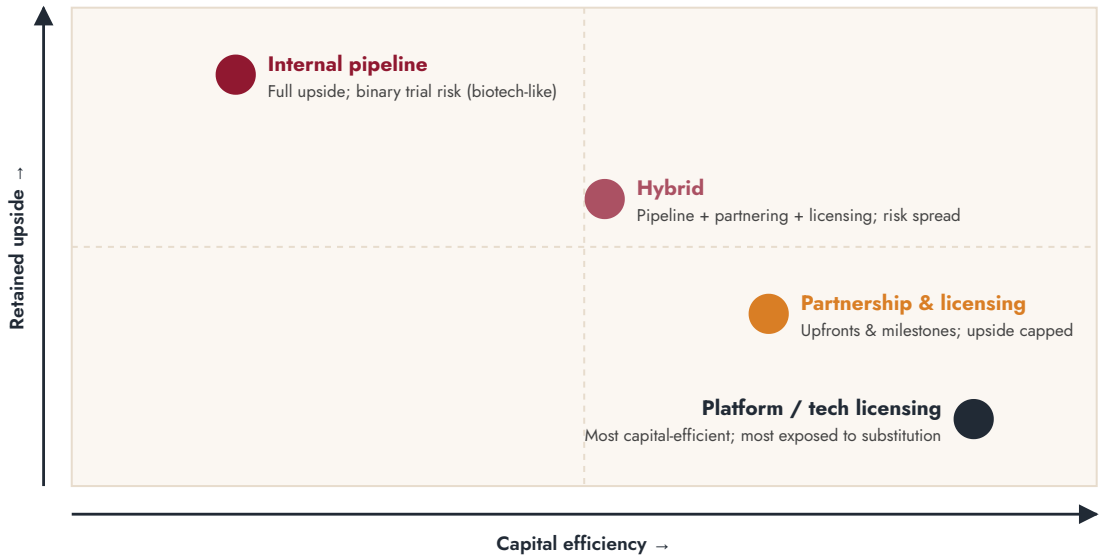

Figure 8 | A qualitative positioning of the commercial models described below. Hybrid models are increasingly the norm.

Source: Vi Partners assessment.

Internal pipeline. The company retains ownership of discovered assets and bears clinical risk directly in exchange for full economic upside, as with Recursion pre-merger or Xaira Therapeutics. This behaves, financially, like a biotech, with binary trial risk, regardless of how computational the discovery process was.

Partnership and licensing. The platform is validated and monetized through upfront and milestone payments from pharma partners who then run development, as with Absci's AstraZeneca, Merck and Almirall deals, or Exscientia's pre-merger agreements with Sanofi and Bristol Myers Squibb. This is more capital-efficient but caps upside, and announced deal sizes are often mostly backend milestones that may never be paid.

Platform or technology licensing. The company sells access to its modeling infrastructure rather than co-owning drug assets, as with Owkin's multi-year platform license to Sanofi or Schrödinger's long-standing software business. This is the most capital-efficient model, but also the one most exposed to substitution: a partner can build the capability in-house, a competitor can undercut it, and openly released models can make a paid tool difficult to justify at all.

Hybrid. Increasingly the norm rather than the exception. Insilico Medicine and Recursion each run proprietary pipeline assets alongside partnered and platform-licensing revenue, spreading risk while retaining some exposure to pipeline upside.

A related pattern is now visible across the sector: platforms that started by selling software access or discovery-as-a-service are increasingly moving toward retaining IP in their own programs and partnering on terms that preserve some exposure to downstream value. What matters in diligence is not only a deal's size but its structure. An agreement where the platform retains meaningful IP, milestones, or royalty exposure behaves very differently, economically, from one where the platform is paid to generate an output that a partner then owns fully. The latter may be a successful services business, but likely it will not receive the valuation characteristics of a scalable platform business.

Two recent events illustrate how this is playing out. Recursion's 2024 acquisition of Exscientia (roughly USD 650–690 million) combined two of the highest-profile AI-native drug discovery companies into a single entity with a deeper pipeline, more partnered programs and a longer cash runway than either had alone. It was a consolidation driven as much by capital efficiency as by strategy. BenevolentAI took the opposite path: having built an internal late-stage pipeline on top of a target-discovery platform, it suffered a costly Phase II failure in 2023 and has since withdrawn to earlier-stage partnering and delisted in March 2025, a reminder that a platform can only carry its own trials if it can survive one failing.

USD 5.1bn
Siemens' 2025 acquisition of Dotmatics, a scientific data and software platform

USD 650–690m
Recursion's 2024 acquisition of Exscientia

Up to USD 1bn
Joint Lilly–NVIDIA co-innovation lab: talent, infrastructure and compute over five years (2026)

Strategic acquirers, meanwhile, are paying up for infrastructure rather than pipeline. Siemens' 2025 acquisition of Dotmatics for USD 5.1 billion valued a scientific data and software platform well above where most AI-native pipeline companies trade, and Eli Lilly's January 2026 co-innovation lab with NVIDIA, with the two companies investing up to USD 1 billion over five years in talent, infrastructure and compute, suggests large pharma increasingly wants to own the computational layer, not merely license access to it.

What determines value?

Computational drug discovery companies are unusual to value because they sit between software, technology platforms, and biotech. Early in their development, investors are primarily funding the scientific insight, computational technology, team, and potential to build a defensible platform. As the platform generates proprietary data and experimentally validated results, value increasingly shifts toward the quality of the learning loop, partnerships, and the assets it can generate.

This creates an important distinction between platform value and pipeline value. A differentiated computational model can be worth a significant value while few competitors can match it, but on its own it rarely holds that value once the methods become widely available. Proprietary data, experimental infrastructure, validated biology, and the ability to repeatedly

generate high-quality therapeutic candidates are what allow platform value to compound. Once candidates enter clinical development, valuation increasingly reflects the probability-adjusted value of the resulting pipeline, alongside the option value of the underlying platform.

This helps explain why the high valuations seen for companies such as Isomorphic Labs and Insilico Medicine should not be interpreted simply as a price for an AI model. The underlying technology matters, but so do proprietary data, validated biology, pipeline optionality, partnerships, and the strategic value of capabilities that are difficult for pharmaceutical companies to build internally. The relative contribution of these factors changes materially with maturity.

The market therefore does not apply a single "AI drug discovery" multiple. The valuation shifts as risk moves: from technology and the range of programs a platform could pursue at the earliest stages, toward validated biology, pipeline potential, and ultimately clinical evidence. Strategic transactions can place additional value on the platform itself, including data, technology, talent, partnerships, and time saved.

This is why we do not evaluate these companies primarily by the sophistication of their models. We look instead at the underlying method, the experimental feedback loop, what the company owns and learns, how it captures economic value, and what evidence exists at its stage. We also weigh how far a company can take a molecule. Discovery is only the first step, and the value increases for those that can follow an asset from bench to bedside or retain exposure when a partner does.

5. What we look for

Excitement about the strength of the underlying science is not a substitute for the diligence questions that determine whether a computational discovery platform can become a durable business. We do not treat the following as five equally weighted boxes to check. Because we invest principally early in the drug development path, we order them by what is observable at that stage rather than by long-run importance. A company should already be able to demonstrate method fit and a working experimental feedback loop, so weakness there counts heavily. Commercial structure and durable proof are largely unbuilt this early, so a thinner answer there might reflect maturity rather than a big flaw.

Figure 9 | The tests are interconnected and are assessed together, not scored and averaged.

1. The right method for the biology. The question is not whether a platform is physics-based or AI-native, but whether the method was chosen because it fits the specific biology being targeted. Physics-based methods can be particularly valuable where mechanistic detail matters and data are limited. Generative and machine-learning methods, which require large training sets, can be powerful where proprietary data exist at meaningful scale. We want a team to explain clearly why its method is suited to its target class and what it can do that a more general approach cannot. We also ask where in the workflow a platform intervenes, and whether that step is one where capacity is genuinely scarce, as with target identification or clinical translation, or one that is already well served, as with molecule generation itself. Equal technical sophistication creates more value when applied to a bottleneck than to a step the industry has already largely solved.

2. Experimental feedback loop. Does the company control the pace of its experimental feedback, or does it depend on someone else's backlog? Owning a wet lab is one way to establish this. Genuine co-development or CRO integration can also work, provided the company still directs experiment design and receives iterative results on a useful timeline. A platform that simply hands a prediction to a partner and waits may still be a good business, but it should not be credited with a closed feedback loop it does not control.

3. Defensible data and learning. Is the data proprietary and difficult to replicate, such as the company's own wet-lab output or exclusive partner data, or is it mainly public data reprocessed with a new model? Public data can support strong science, but it is harder to defend as a business unless the company adds proprietary information or capabilities. We also ask whether the platform's predictions have been confirmed experimentally, rather than only benchmarked retrospectively against public test sets. The important question is whether each experimental cycle makes the platform more difficult to replicate.

4. Platform rather than service economics. Does the platform work across multiple targets, modalities, or indications with the same core technology, or does every new program require materially new science? A useful test is: if the company stopped development today, what could it still sell tomorrow? If each new engagement compounds the underlying models, data, or capabilities, the business is becoming a platform. If each engagement produces another bespoke analysis, it is a service business regardless of how much AI sits underneath it. The commercial model should also match the capital the company has to absorb a failure: a platform that runs its own programs needs to be funded well enough to survive one of them failing, while a partnering or licensing model does not.

5. Evidence appropriate to the stage. What counts as proof changes with a company's maturity, but the underlying test does not: has the company produced any real, prospectively confirmed result yet? Because we invest principally early, the proof required at a later stage does not exist yet, so the bar should be the smallest verifiable result a company of that age could plausibly have produced. At this stage, that may mean closing a design-test-learn loop on a real target with experimental confirmation. Later rounds should raise the bar toward paid milestones, repeat partnerships, and ultimately clinical evidence. What we want to avoid at any stage is a platform whose evidence remains entirely retrospective or simulated.

A short list, not a scorecard

A company that clears all five tests is genuinely uncommon, which is the point: this is intended to produce a short list, not a scorecard everyone passes with a good enough average. The tests are also interconnected. Strong science without a feedback loop is difficult to compound. Proprietary data without method fit is difficult to translate, and attractive partnerships without retained economic value do not necessarily create a durable platform. What we look for at the stage we invest is credible evidence that the company can turn its scientific approach into a learning system and, ultimately, into economic value.

6. Market size, funding and clinical evidence

Market size estimates for AI and computational drug discovery vary widely depending on how broadly the category is defined, and the estimates themselves continue to evolve. Even for the same base year, published estimates differ by a factor of more than three. MarketsandMarkets values the AI-in-drug-discovery market at USD 3.92 billion in 2025 and projects USD 17.56 billion by 2031, a 28.1 percent CAGR over 2026–2031.[12] Grand View Research puts it at USD 2.3 billion in 2025, growing to USD 13.8 billion by 2033 at a 24.8 percent CAGR over 2026–2033.[13] A 2026 report from Research and Markets is considerably more conservative: USD 1.13 billion in 2025, growing to USD 2.29 billion by 2031, a CAGR of about 12.5 percent.[14] The range is a reminder that market estimates depend heavily on category boundaries and adoption assumptions. The more durable signal is the increasing willingness of pharmaceutical companies to commit capital to computational discovery infrastructure and partnerships.

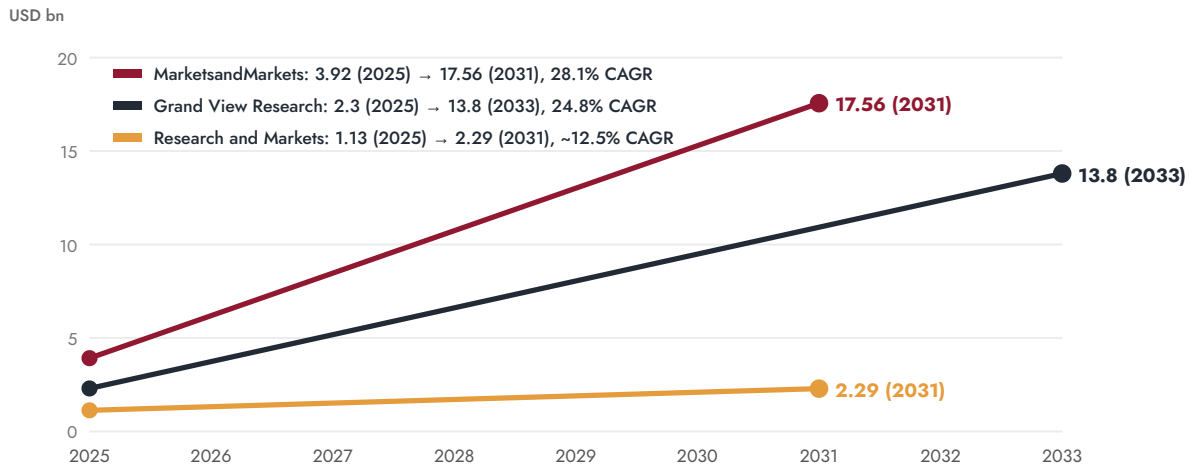

Figure 10 | Three published forecasts, start and end points as reported. All three project growth, but their 2025 estimates differ by more than a factor of three, underscoring how much they depend on category definitions and assumptions.

Sources: MarketsandMarkets (July 2026); Grand View Research (2026, updated June 2026); Research and Markets via GlobeNewswire (May 2026).

Venture funding has kept pace regardless. 2024 and 2025 combined saw roughly USD 19.9 billion invested across more than 600 rounds in the sector, with 2025 funding of around USD 11 billion ahead of 2024's USD 8.9 billion, despite a broader biotech funding slowdown elsewhere.[15] Big pharma has moved from pilot projects to structural commitments: combined deal value between Isomorphic Labs and its Eli Lilly and Novartis partnerships now approaches USD 3 billion, alongside a separate, newer collaboration with Johnson & Johnson. Sanofi alone has signed platform deals with both Exscientia and Owkin. AstraZeneca, Merck and Almirall have each signed partnerships with Absci.

USD 19.9bn
Venture funding in AI drug discovery, 2024–2025 combined

600+
Financing rounds over the same period

USD 11bn
2025 funding, ahead of USD 8.9bn in 2024

~USD 3bn
Isomorphic Labs' combined potential Lilly and Novartis deal value

The clinical evidence, while still early, is more encouraging than the market-sizing debate might suggest. An independent analysis of AI-discovered molecules in the clinic found Phase I success rates of 80 to 90 percent, well above historical industry averages of roughly 40 to 65 percent, a real signal that AI-assisted design is improving early safety and ADMET

(absorption, distribution, metabolism, excretion, toxicity) prediction.[16,17] Phase II success for the same molecules, however, reverted to around 40 percent, statistically indistinguishable from conventional drug development.

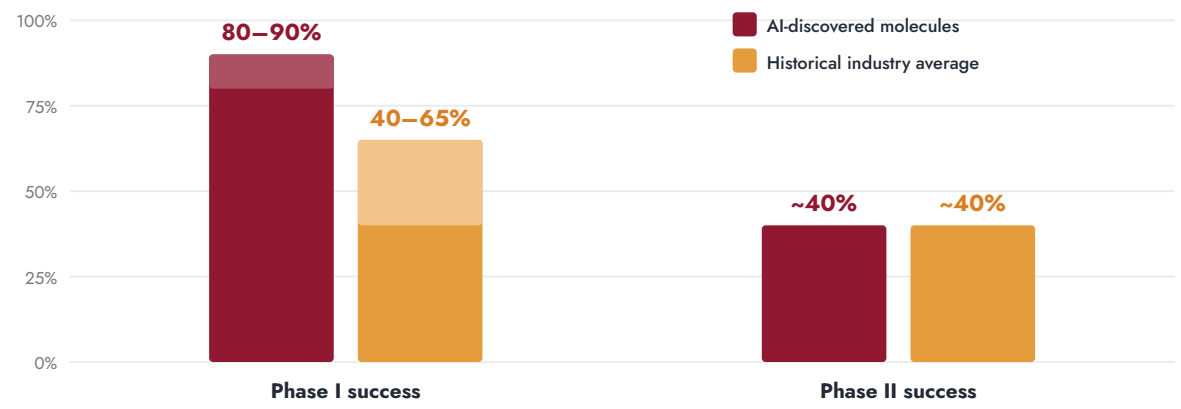

Figure 11 | Computational methods have so far improved the odds of getting a safe molecule into the clinic more than the odds of it working in patients.

Source: Jayatunga et al., Drug Discovery Today (2024); BCG (2025).

> Combined across phases, the analysis put the overall probability of a molecule advancing through clinical development at close to double the historical base rate.[16] This is meaningful, but a reminder that computational methods have so far changed the odds of getting a safe molecule into the clinic more than the odds of it eventually working in patients.

This split is consistent with where the computational methods in this paper actually intervene. Optimizing a molecule's potency, selectivity or ADMET profile (the work most generative and physics-based platforms are built for) directly affects whether it clears Phase I safely. It has much less to say about whether the target itself is the right one in humans, or whether the trial around it is built to detect a real effect in the right patients: closer to what AstraZeneca's own well-known review of its pipeline called the right target and the right patient, still among the clearest diagnoses of why efficacy fails.[18] The categories most likely to move Phase II and III over time are the ones working further upstream and downstream of molecule design: target discovery platforms that use causal inference and patient-derived multi-omic data to test whether a target actually drives disease in humans, rather than in a mouse or a cell line (Section 3.1), and digital biology platforms building more human-relevant cell and tissue models to test hypotheses before the clinic (Section 3.5).

Digital twins: the trial-side complement

On the trial side, digital twins (patient-specific computational models used to design trials, build synthetic control arms, and match patients to the regimen most likely to work for them) are a complementary technology we covered in depth in our earlier white paper, *Digital Twins in Oncology*,[19] and one of the clearer paths to improving efficacy read-outs once a molecule already exists. Digital twins address trial design and analysis, principally who is enrolled in which arm and how a signal is read out, rather than the execution and relationship work of recruiting, screening and retaining patients in the first place, which remains a largely human, not computational, constraint.

A parallel shift is reshaping who is building these platforms in the first place. The large horizontal AI labs no longer treat selling tools into biology from the outside as a sufficient position, and several are now building or buying their way in directly. Google's Isomorphic Labs is the most advanced example (Section 3), but it is no longer the only one. Anthropic reportedly acquired Coefficient Bio, a small AI drug discovery startup, for roughly USD 400 million in April 2026, and in June 2026 launched Claude Science, a product built to support scientific research, including computational biology and drug development, which Anthropic also plans to use for its own research into drugs for rare, neglected diseases.[20–22] OpenAI has taken a software-first approach with GPT-Rosalind, a reasoning model for biology and drug discovery now used by Amgen, Moderna and Thermo Fisher Scientific. Amazon Web Services launched Amazon Bio Discovery, pairing foundation models with CRO partners such as Twist Bioscience and Ginkgo Bioworks for wet-lab validation.[20] NVIDIA has moved from supplying compute to co-running discovery itself, including a dedicated lab with Eli Lilly announced in January 2026.[23]

Horizontal AI labs are moving into biology directly

Company	Vehicle	Approach
Google	Isomorphic Labs	Spin-out building its own pipeline; USD 2.1bn Series B (May 2026) and ~USD 3bn in partnerships
Anthropic	Coefficient Bio; Claude Science	Reported acquisition (~USD 400m, April 2026); research product; own research on drugs for rare, neglected diseases
OpenAI	GPT-Rosalind	Reasoning model for biology used by Amgen, Moderna, Thermo Fisher
Amazon (AWS)	Amazon Bio Discovery	Foundation models paired with CRO partners for wet-lab validation
NVIDIA	Co-innovation lab with Eli Lilly	From supplying compute to co-running discovery (January 2026)

Figure 12 | Several general-purpose AI players are building or buying their way into discovery rather than only selling tools into it.

Sources: GEN (May 2026); Fierce Biotech (April 2026); MIT Technology Review (June 2026); NVIDIA (January 2026).

Consolidation is also underway on the biotech side. Beyond Recursion's acquisition of Exscientia, 2025-2026 saw repeated smaller M&A activity across AI-native biotech, pharma, and data and CRO players, alongside acquisitions of infrastructure providers such as Dotmatics. Those valuations suggest the market is beginning to price computational drug discovery infrastructure as durable technology rather than speculative biotech. Both the technology and the capital are moving toward clinical programs, even if most of those programs have yet to read out: Insilico Medicine reached that point in July 2026, when its internally discovered rentosertib became one of the first AI-originated drug candidates to enter Phase III trials.[24]

7. How we see the opportunity

We think the categories described in this paper will keep converging rather than staying separate, and that the winners over the next years will be defined less by which specific algorithm they use and more by how tightly they close the loop between computation and experiment, between what a model predicts and what a real molecule does, and how disciplined they are about the commercial model they choose for the risk they are actually equipped to carry. We assess these qualities through the five parameters set out in Section 5, which are designed to be applied before clinical evidence exists.

PORTFOLIO SPOTLIGHT · SIBYLLA BIOTECH

Our own portfolio gives us a close-up view of what these tests can look like at an early stage. Sibylla Biotech is a particularly clear example of a method matched to its biology: its ability to target folding intermediates rather than final protein structures makes it possible to address protein classes, such as KRAS and cyclin D1, that have resisted conventional structure-based design for decades. Its collaborations with Takeda and MD Anderson Cancer Center also provide the kind of early external validation that can be meaningful before clinical proof exists.

More broadly, the landscape reinforces our view that there is no single computational drug-discovery category: different biological problems call for different combinations of data, modeling, experimentation, and commercial structure.

We look for healthcare technologies that reduce the cost and time of getting an effective therapy to the right patient. Computational drug discovery addresses both. **We believe the strongest opportunities will be built by companies that close the loop between biological insight, computation, and experiment, while retaining enough proprietary data, IP, or economic participation that each experimental cycle makes the platform harder to replicate.**

8. Outlook

Drug discovery will keep becoming more computational, but not because computation replaces the wet lab. Every platform discussed in this paper still depends on real experiments to validate its predictions. What is changing is which experiments get run, in what order, and how much is learned from each one before the next is designed. Machine learning narrows the search space, generative models propose candidates that would not have been found by screening alone, and physics-based methods provide mechanistic grounding when data run out. Where these approaches converge, discovery becomes less of a linear funnel and more of a learning system: computation proposes, experiments test, and each cycle improves what comes next.

> **The competitive advantage will increasingly come from how quickly a company can move around that loop, how much proprietary information it accumulates, and how effectively it converts that learning into assets or recurring economic value.**

One structural risk sits above every category described in this paper. As Section 6 describes, the large horizontal AI labs are moving into biology directly rather than supplying tools to it. If their general-purpose models keep improving, a company whose only advantage is a biology-specific model or dataset may see that advantage matched sooner than it expects. For specialized computational discovery platforms, that makes timing, defensibility, experimental integration, and speed of execution increasingly important. This does not make the category less investable, but it does raise the standard for what a company must own that a general-purpose model cannot replicate.

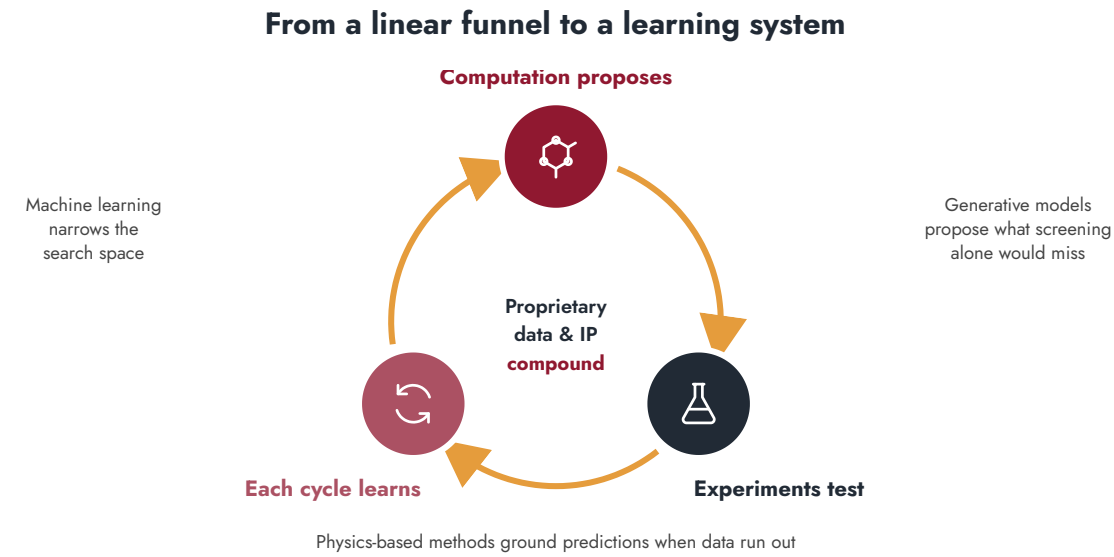

Figure 13 | Competitive advantage increasingly comes from the speed of the loop and what each cycle leaves behind.

None of this makes the winners easy to identify in advance. The five tests in Section 5 are not a scoring system that leads to a single answer. They are separate dimensions that have to be assessed together, because a company can look strong on any one of them while failing on another in a way that cancels out the strength. A well-chosen method can sit on an experimental cycle too slow to improve it, and genuinely proprietary data can sit under a commercial model that captures none of its value. At seed and Series A there is not yet enough history to confirm that a platform compounds, so the judgment rests on whether the method, the experimental cycle, and what the company owns are set up to make that happen.

References

1. Insilico Medicine. Insilico and Lilly enter a research and licensing collaboration to advance AI-driven drug discovery. Press release (November 10, 2025). https://insilico.com/news/e9xm9y0if1-insilico-and-lilly-enter-a-research-and
2. Scannell, J.W., Blanckley, A., Boldon, H., & Warrington, B. Diagnosing the decline in pharmaceutical R&D efficiency. Nature Reviews Drug Discovery 11, 191–200 (2012). https://www.nature.com/articles/nrd3681
3. DiMasi, J.A., Grabowski, H.G., & Hansen, R.W. Innovation in the pharmaceutical industry: New estimates of R&D costs. Journal of Health Economics 47, 20–33 (2016). https://doi.org/10.1016/j.jhealeco.2016.01.012
4. Deloitte. Measuring the return from pharmaceutical innovation, 2023–2025 editions. https://www.deloitte.com/ch/en/Industries/life-sciences-health-care/research/measuring-return-from-pharmaceutical-innovation.html
5. BIO, Informa Pharma Intelligence & QLS Advisors. Clinical Development Success Rates and Contributing Factors, 2011–2020 (2021). https://www.bio.org/press-release/bio-releases-largest-study-ever-clinical-development-success-rates
6. Leenaars, C.H.C., Kouwenaar, C., Stafleu, F.R., Bleich, A., Ritskes-Hoitinga, M., De Vries, R.B.M., & Meijboom, F.L.B. Animal to human translation: a systematic scoping review of reported concordance rates. Journal of Translational Medicine 17, 223 (2019). https://doi.org/10.1186/s12967-019-1976-2
7. Harrison, R.K. Phase II and phase III failures: 2013–2015. Nature Reviews Drug Discovery 15, 817–818 (2016). https://www.nature.com/articles/nrd.2016.184
8. Insilico Medicine. From Start to Phase 1 in 30 Months. Press release (2022). https://insilico.com/phase1
9. Ferruz, N., Schmidt, S., & Höcker, B. ProtGPT2 is a deep unsupervised language model for protein design. Nature Communications 13, 4348 (2022). https://doi.org/10.1038/s41467-022-32007-7
10. Ross, G.A., et al. The maximal and current accuracy of rigorous protein-ligand binding free energy calculations. Communications Chemistry 6, 222 (2023). https://www.nature.com/articles/s42004-023-01019-9
11. The Nobel Prize in Chemistry 2024. Press release, Nobel Foundation (2024). https://www.nobelprize.org/prizes/chemistry/2024/press-release/
12. MarketsandMarkets. AI in Drug Discovery Market – Global Forecast to 2031. Published July 2026. https://www.marketsandmarkets.com/Market-Reports/ai-in-drug-discovery-market-151193446.html
13. Grand View Research. Artificial Intelligence in Drug Discovery Market Size, Share & Trends Analysis Report, 2026–2033. Report ID GVR-4-68039-029-7 (2026; last updated June 2026). https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-drug-discovery-market
14. Research and Markets. AI in Drug Discovery Research Report 2026: Global Market Analysis, Competitive Landscape, Opportunities, and Forecasts, 2021–2025 & 2025–2031. Distributed via GlobeNewswire (May 8, 2026). https://www.globenewswire.com/news-release/2026/05/08/3290907/28124/en/AI-in-Drug-Discovery-Research-Report-2026-Global-Market-Analysis-Competitive-Landscape-Opportunities-and-Forecasts-2021-2025-2025-2031.html
15. DealForma. AI-ML Drug Discovery and Licensing: R&D, M&A, Ventures and IPOs – 2025 Review (January 13, 2026). https://dealforma.com/ai-ml-drug-discovery-and-licensing-rd-ma-ventures-and-ipos-2025-review/
16. Jayatunga, M.K.P., et al. How successful are AI-discovered drugs in clinical trials? A first analysis and emerging lessons. Drug Discovery Today 29(6), 104009 (2024). https://doi.org/10.1016/j.drudis.2024.104009
17. Boston Consulting Group. Reigniting Biopharma's Research Engine (2025). https://www.bcg.com/publications/2025/reigniting-biopharmas-research-engine
18. Cook, D., Brown, D., Alexander, R., March, R., Morgan, P., Satterthwaite, G., & Pangalos, M.N. Lessons learned from the fate of AstraZeneca's drug pipeline: a five-dimensional framework. Nature Reviews Drug Discovery 13, 419–431 (2014). https://www.nature.com/articles/nrd4309
19. Escala-Garcia, M. Digital Twins in Oncology: Seizing Opportunities Across Clinical Care, Drug Development, and Health Systems. Vi Partners white paper (January 2026). https://a.storyblok.com/f/130367/x/898f86516e/digital-twins-in-oncology.pdf
20. Lin, F. Big Tech Targets Drug Discovery with Wave of Life Science Platforms. GEN – Genetic Engineering and Biotechnology News (May 4, 2026). https://www.genengnews.com/topics/artificial-intelligence/big-tech-targets-drug-discovery-with-wave-of-life-science-platforms/
21. Maddox, W. Anthropic acquires stealth AI startup Coefficient Bio in $400M deal: reports. Fierce Biotech (April 6, 2026). https://www.fiercebiotech.com/biotech/anthropic-acquires-stealth-ai-startup-coefficient-bio-400m-deal
22. Huckins, G. Claude Science is Anthropic's newest flagship product. MIT Technology Review (June 30, 2026). https://www.technologyreview.com/2026/06/30/1139987/claude-science-is-anthropics-newest-flagship-product/
23. NVIDIA Corporation and Eli Lilly and Company. NVIDIA and Lilly Announce Co-Innovation AI Lab to Reinvent Drug Discovery in the Age of AI. Press release (January 12, 2026). https://nvidianews.nvidia.com/news/nvidia-and-lilly-announce-co-innovation-lab-to-reinvent-drug-discovery-in-the-age-of-ai
24. Insilico Medicine. Insilico Initiates Phase III Clinical Trial of Rentosertib. Press release (July 7, 2026). https://insilico.com/news/xmjsn4l091-insilico-initiates-phase-iii-clinical-tr

About the Author

Maria Escala-Garcia is an Investment Associate at Vi Partners, where she focuses on healthcare and life sciences investments, with particular interest in data-driven innovation, digital health, and computational approaches to medicine.

She holds a PhD in bioinformatics and genetic epidemiology and has an academic and business background spanning computational oncology, precision medicine, and digital health. Prior to joining Vi Partners, she worked across the healthtech and consulting sectors, supporting the development, evaluation, and adoption of data-enabled healthcare technologies at the interface of science, clinical practice, and innovation.

Disclaimer

This white paper reflects the author's views and is for informational purposes only. It does not constitute investment advice or an offer to invest. The statements regarding market size, growth rates, funding amounts, and company positioning are based on published reports and public information as of October 2026 and are subject to change.