AI Drug Discovery: Perceptic's Unified Pharma R&D Platform

5/12/2026 • 45 min read

ai drug discovery | pharmaceutical r&d | perceptic | machine learning in pharma | biotech data integration | drug development lifecycle | clinical data analysis

Executive Summary

Perceptic, a London- and Basel-based startup founded in 2024 by three former Palantir Life Sciences executives, has emerged from stealth with a $12 million seed round led by Accel (with participation from Air Street Capital and Elder Gull) ([1] fortune.com) ([2] techfundingnews.com). Its mission is to provide the pharmaceutical industry with an **"AI operating system for drug development"** – a unified platform that spans the entire drug discovery and development lifecycle ([3] fortune.com) ([4] www.perceptic.com). Unlike most AI drug discovery startups that focus on narrow slices of R&D, Perceptic aims to act as the **"connective tissue"** linking disparate AI tools and data sources, so that insights from target identification to clinical trial design compound rather than "dying at each handoff" ([5] fortune.com) ([6] techfundingnews.com). The founders – Tilman Flock (CEO), Zaki Trache, and Martin Copes – leveraged their experience building Palantir's AI platform and Life Sciences practice to create an infrastructure- and model-agnostic system. Firms like CSL (Australia) are already deploying Perceptic's platform to accelerate asset scouting, indication selection, and clinical data analysis ([7] fortune.com) ([8] european-biotechnology.com).

Perceptic's emergence comes amid a surge of AI-driven drug discovery initiatives. The global *AI in drug discovery* market was about **$1.72 billion in 2024** and is projected to exceed **$8.5 billion by 2030** (a CAGR >30%) ([9] techfundingnews.com). Venture capital has heavily funded the space – for example, Google/DeepMind's spinout Isomorphic Labs raised a **$2.1 billion Series B** in May 2026 ([10] techfundingnews.com), and Insilico Medicine recently closed a **$2.75 billion preclinical portfolio deal with Eli Lilly** ([11] techfundingnews.com). Dozens of startups (Recursion, Exscientia, FormationBio, Xaira, RedwoodBio, and many others) are developing AI tools for individual R&D stages. However, no novel drug wholly discovered by AI has yet completed clinical trials, and industry experts note that simply cranking up "point solutions" is unlikely to realize the full potential of AI ([12] fortune.com) ([13] moneyweek.com). Instead, observers argue the **next frontier is an integrated platform** – an "operating system" – that harmonizes data and decisions across the 10–15 year drug-development cycle ([14] fortune.com) ([6] techfundingnews.com).

This report provides a comprehensive analysis of Perceptic's approach and its context in the evolving AI-based R&D landscape:

- It starts with background on the massive cost, complexity, and data fragmentation in traditional pharmaceutical R&D.
- It surveys recent advances and setbacks in AI-driven drug discovery, citing case studies and market data.
- It examines the concept of a **"data foundation" for R&D**, including knowledge graphs and unified analytics, and compares Perceptic's strategy to other approaches (e.g. IntuitionLabs, Ontotext, Stardog, BenevolentAI).
- It details Perceptic's technology (its "Scout", "PercepticOS", and "Atlas" modules), founders' vision, early results, and investor rationale, drawing on statements by the founders, investors (Accel's Sonali De Rycker, Air Street's Nathan Benaich), and external analysts ([15] fortune.com) ([16] techfundingnews.com) ([17] techfundingnews.com).
- It integrates quantitative evidence (market forecasts, AI R&D success rates, performance improvements) and qualitative viewpoints (executive interviews, academic & industry commentary) to provide a balanced view.
- Finally, it discusses the broader implications for the pharma industry and future directions, including adoption challenges, regulatory considerations, and potential impact on healthcare and global innovation.

Key Findings: Perceptic is positioning itself to fill a critical gap by linking siloed datasets and AI tools into a cohesive workflow. In early deployments, its platform has dramatically accelerated tasks like asset screening (from hundreds per week to thousands in minutes) and due diligence (compressing weeks of analysis into hours), while improving data traceability ([18] european-biotechnology.com) ([19] press.airstreet.com). Whether the industry ultimately consolidates around an "R&D OS" remains to be seen, but many investors and insiders believe such platforms will be the catalyst for the next wave of biopharma innovation ([14] fortune.com) ([20] techfundingnews.com).

Background: The Challenges of Pharmaceutical R&D

The Traditional R&D Paradigm

Developing a new drug is notoriously **long, costly, and risky**. A 2023 analysis reported that bringing a medicine from initial concept to market often takes **more than a decade** and can exceed **$2 billion** in total cost ([21] techfundingnews.com). Average success rates are low: only a small fraction of candidate molecules survive preclinical screening, clinical trials are costly and often fail (traditional phase I success rates are on the order of 40–65% ([22] www.deloitte.com)), and even approved drugs may only address limited patient populations.These challenges persist despite enormous R&D expenditures worldwide.

Several structural factors contribute to this inefficiency:

- **Fragmented workflows:** The drug lifecycle spans target discovery, molecular design, preclinical testing, clinical trials (Phase I-III), regulatory submission, and post-market analysis. Each stage traditionally involves different groups, tools, and data systems. When moving from one phase to another, vital insights can be lost. As one venture investor notes, “we all know [drug development] took more than a decade to reach patients… and it also ran on fragmented systems that barely communicate with each other – research insights disappearing between teams, clinical data in silos, and billion-dollar decisions depending on manually stitching incomplete information” ([21] techfundingnews.com). In this linear handoff model, any improvements in one silo often do not propagate, and as Perceptic co-founder Tilman Flock observes, “insight dies at every handoff” unless kept connected ([5] fortune.com).
- **Complex, heterogeneous data:** Pharmaceutical R&D generates massive and diverse data: genomic and proteomic databases, scientific literature, experimental assays, preclinical animal studies, electronic health records, clinical trial datasets, and proprietary corporate research archives. Each data type (public databases like PubMed or clinical trial registries, internally generated omics data, commercial datasets from vendors, etc.) uses different formats and standards ([23] fortune.com). Integrating these sources manually is extremely difficult and error-prone.
- **High regulatory standards:** Any findings or workflows that influence pandemic must be fully auditable. Data provenance and traceability are paramount, meaning automated tools cannot “hallucinate” or fabricate knowledge without clear sourcing. As Flock notes, “customers cannot tolerate AI hallucinations… [so] our system allows customers to trace every claim back to its source” ([24] fortune.com). This requirement raises the bar for AI solutions beyond typical e-commerce or media use cases.
- **Risk-averse culture:** Pharmaceutical companies have traditionally been cautious about new technology, due to lives being at stake. Established giants often have entrenched departmental structures (biology vs chemistry vs clinical development, etc.), legacy IT systems, and long product cycles that make radical process changes slow.

These factors have meant that, despite decades of digital innovation in areas like finance and manufacturing, pharma R&D has lagged in productivity gains. However, the convergence of large biomedical datasets, powerful computing (e.g. cloud and GPU clusters), and breakthroughs in machine learning is beginning to offer new opportunities to **re-shape** how drugs are discovered and developed ([25] www.deloitte.com) ([26] www.axios.com).

The AI Hype and Reality in Drug Discovery

The notion that AI could **revolutionize drug discovery** has been widely touted. Pharmaceutical executives and technologists alike envision a future where AI-driven models dramatically compress timelines, identify far more promising candidates, and even suggest novel therapies. A full set of “frontend” AI models could, in theory, design molecules that

precisely target disease-relevant proteins, optimize biochemical properties, and predict safety profiles ([27] www.deloitte.com). On the clinical side, machine learning and automation could streamline patient recruitment, optimize trial protocols, and synthesize real-world evidence to validate efficacy ([27] www.deloitte.com) ([28] time.com).

Some high-profile AI accomplishments fuel this optimism. For instance, Google DeepMind's **AlphaFold 2** (2021) demonstrated the ability to predict protein 3D structures from amino acid sequences with unprecedented accuracy, potentially unlocking targets that were previously enigmatic ([29] moneyweek.com). Other examples include:

- **Generative chemistry:** Companies like **Insilico Medicine** and **Exscientia** use deep learning to suggest new molecular structures expected to bind a given protein target. Insilico recently secured a $2.75 billion preclinical portfolio deal with Eli Lilly, underscoring confidence in this approach ([11] techfundingnews.com).
- **Phenotypic screening with AI: Recursion Pharmaceuticals** has amassed billions of cell imagery data and uses AI to identify phenotypic patterns across disease models, reportedly discovering novel therapeutic leads in oncology and immunology (it has partnerships with Bayer and Roche).
- **Clinical AI:** Startups like **Formation Bio** apply AI to optimize clinical trial diagnostics and logistics. Formation Bio CEO Ben Liu reports that their platform can **"save as much as 50% of the time of a trial"** by automating administrative tasks (recruitment, filings, etc.) ([30] time.com), and already has sold two AI-accelerated candidates to major pharma (Sanofi and Lilly) for nearly $2 billion.

At the same time, industry observers caution that **actual impact has been slower than hoped**. Over the past decade, despite large investments in AI for drug R&D, the annual number of new FDA drug approvals has held broadly at 40–60 per year, roughly unchanged – leading critics to question whether AI is truly delivering a breakthrough ([12] fortune.com) ([31] time.com). (Notably, however, Deloitte reports that as of late 2023, AI-discovered molecules had achieved an exceptionally high Phase I success rate (80–90%), though few have reached end-to-end approval ([22] www.deloitte.com).) Some contributing reasons include incomplete biological knowledge ("we're trying to intervene in a system we only slightly understand," notes Insitro's Daphne Koller ([32] apnews.com)), and the intrinsic complexity of diseases. Many leading companies – GSK, Novartis, Bayer, Roche, etc. – now embed AI into certain programs, but often as one tool among many ([33] moneyweek.com) ([26] www.axios.com).

As one analysis puts it, **"AI has been billed for years as an avenue to medical advancements, whether through drug discovery, diagnostics or treatment – but that moment may be disruptive [for pharma]"** ([34] www.axios.com). The key challenge is not just having better models, but integrating them into the deeply siloed, data-rich workflows of biopharma. This has led venture investors to emphasize the concept of an **AI platform or "operating system" for pharma R&D**. Instead of selling isolated model libraries (e.g. for molecule generation or patient matching), these new solutions aim to tie together data, models, and decisions across the entire R&D process ([14] fortune.com) ([6] techfundingnews.com). The hypothesis is that significant compounding benefits arise when insights accumulate rather than dissipate between stages.

The Data Foundation Race: Unifying R&D Data

Underpinning the AI revolution in R&D is the idea of a robust **data and knowledge foundation**. Modern AI models thrive on large, well-organized datasets. Pharma companies, however, typically have vast amounts of legacy data scattered across systems. To leverage AI effectively, many firms and startups seek to build a unified data layer—a single, trusted source of truth—covering the breadth of biomedical knowledge.

Knowledge Graphs and Semantic Integration. A common approach is the use of *knowledge graphs*: structured networks where concepts (genes, diseases, trials, compounds, etc.) are nodes and relationships are edges. By mapping diverse datasets into a shared graph ontology, these platforms enable rich querying and inference. For example, Ontotext's "Knowledge Graph Drug Discovery" solution ingests public databases (PubMed, UniProt, NCBI, patents, etc.) alongside a company's proprietary data, normalizes all information to common ontologies, and uses semantic reasoning

to uncover hidden correlations◆ ([35] www.ontotext.com) ([36] www.ontotext.com). Similarly, Ontoforce's DISQOVER platform links siloed data for pharma organizations, promising faster searches and enhanced decision-making ([37] www.ontoforce.com) ([38] www.ontoforce.com). Even broader enterprise AI platforms (e.g. Palantir Foundry with its AIP component) now offer life-sciences-specific modules that ingest and harmonize clinical trial records, lab data, regulatory documents, and more into a unified environment ([39] www.palantir.com) .

These knowledge-graph or unified-data solutions typically advertise benefits such as **"FAIR" data** (Findable, Accessible, Interoperable, Reusable), improved traceability, and the ability for ML models to train on enriched, linked datasets ([40] www.stardog.com) ([41] www.stardog.com). For instance, Stardog's drug discovery platform uses virtualization to connect all R&D data (across labs and formats) so scientists can "follow the chain of custody and talk about the 'why' of the data," according to a biotech user ([42] www.stardog.com) ([43] www.stardog.com). In practice, these efforts can dramatically speed tasks: Stardog notes that unified data lets a researcher quickly find all similar compounds tested or assays already run for a target, whereas without it they might duplicate experiments.

GenAI and Workflow Orchestration. Beyond knowledge graphs, a new trend is leveraging generative AI (LLMs, agents) to create *operational layers* on top of the data. For example, one startup (IntuitionLabs) is specializing in generative AI solutions integrated with pharma tech stacks (Veeva, CRM, etc.), aiming to automate tasks like report generation or study management. Domain-specific AI "agents" can be tuned to particular R&D sub-tasks: a molecule-design LLM, a clinical-text summarizer, a real-world evidence processor, etc. The challenge is connecting these agents so that their outputs feed into one another coherently — the very "OS" problem that Perceptic addresses.

The Race Heats Up. The notion of an R&D data platform has attracted many players. Big tech companies are exploring it: in April 2024, Nvidia announced projects to build specialized generative models and supercomputers for pharma (in partnership with Lilly, among others), explicitly targeting end-to-end drug R&D workflows ([44] www.axios.com) ([45] www.axios.com). Established R&D vendors (like Veeva, LabCorp, IQVIA) are also adding AI layers to their offerings. Venture-backed startups like BenevolentAI and Embecta have worked on their own knowledge graphs and mining tools for years. Recently, a cluster of EU-funded "HealthTech" initiatives and even national AI centers (e.g., Boehringer Ingelheim's AI hub in London) demonstrate growing ecosystem support ([46] european-biotechnology.com). This can be viewed as a *race* to lock down proprietary data and build group-specific AI pipelines before others.

In summary, the **data foundation race** is about overcoming the fragmentation in pharma R&D data. Companies that succeed in creating a comprehensive, curated data platform – one that can ingest internal records, published science, and commercial datasets with full traceability – will set the stage for powerful AI analytics and automation across discovery, preclinical, and clinical research.

Perceptic: An End-to-End AI Platform

Founders and Vision

Perceptic was co-founded by **Tilman Flock, Martin Copes, and Zaki Trache**, all of whom were senior engineers at Palantir who helped build Palantir's AIP (Artificial Intelligence Platform) and led its Life Sciences practice ([2] techfundingnews.com) ([47] european-biotechnology.com). The founders, drawing on Palantir's rigor around security and data integration, identified what they saw as a critical gap: the pharma industry lacked an integrated AI solution that could "follow the drug" through all R&D stages. They decided to "solve the problem worth solving" of connecting the fragmented workflows and data silos ([48] techfundingnews.com) ([47] european-biotechnology.com). In 2024 they incorporated Perceptic (headquartered in London, with a presence also in Basel) and spent the next two years working quietly with industry partners to design an AI system tailored to pharma's needs ([47] european-biotechnology.com) ([49] press.airstreet.com).

Their vision, as articulated in investor and press communications, is to provide **one intelligence layer across the drug lifecycle** ([50] www.perceptic.com). Perceptic's landing page describes an "AI operating system" that connects evidence,

data, and workflows so that “every insight compounds” – meaning discoveries in one area automatically inform others ([50] www.perceptic.com) ([51] www.perceptic.com). Unlike point solutions that might predict protein structures or screen molecules in isolation, Perceptic is designed to be **infrastructure- and model- agnostic** ([15] fortune.com). In other words, customers can plug in their own hardware resources, legacy data, existing ML models, and Perceptic will tie all these elements together transparently ([15] fortune.com). This approach was deliberate: as Accel’s Sonali De Rycker notes, the goal is to unify *all* the sources of knowledge (internal records, scientific literature, clinical data, competitive intelligence) on one platform rather than improving “each part of the process separately” ([5] fortune.com) ([23] fortune.com).

Importantly, Perceptic’s founders emphasize **traceability and trust**. They recognize that pharma customers demand provenance for any AI-derived conclusion. This is reflected in the platform’s design: all “claims” or recommendations made by the AI can be traced back to source data (papers, trial reports, internal files, etc.), completely avoiding the “hallucinations” afflicting general LLMs ([24] fortune.com). In practice, this means every AI-driven insight has an audit trail. In marketing terms, Perceptic calls its system “connective tissue” – extending existing scientific infrastructure and weaving in AI workers (or agents) that continuously learn from the organization’s accumulated knowledge ([52] www.perceptic.com) ([51] www.perceptic.com).

The founding team is small – as of May 2026 around 15–20 people – but deep in experience. All three have spent years deploying Palantir systems in complex enterprise contexts; now they apply that discipline to drug R&D ([53] www.perceptic.com). Accel, which had been tracking this team since their Palantir days, decided to lead a seed round after seeing prototype use cases in “paid production” with partner companies ([54] fortune.com) ([55] techfundingnews.com). The combination of their creation story and investor endorsements (Accel, Air Street, Elder Gull) is part of what pitches Perceptic as having a “right to win” in Europe, given the concentration of pharma talent in the UK and Switzerland ([56] fortune.com) ([46] european-biotechnology.com).

Platform Architecture: Scout, PercepticOS, Atlas

Perceptic’s platform is structured around three core AI **applications** – internally named **Scout**, **PercepticOS**, and **Atlas** ([19] press.airstreet.com) ([57] press.airstreet.com) – all running on a unified back-end (the “AI intelligence layer”). Each addresses a different aspect of the R&D workflow:

- **Scout (Asset Scouting Engine)**: Scans and triages external drug assets (licensing candidates, competitors’ pipelines, academic patents, etc.) relevant to a customer’s research goals. In practice, a scientist might input a hypothesis or therapeutic target into the system, and Scout will use AI to evaluate thousands of external peptide libraries, small molecules, and published programs, ranking them against the hypothesis ([58] press.airstreet.com) ([59] european-biotechnology.com). In live usage, Perceptic reports that **evaluation time has dropped from weeks to 1 hour**, and asset screening throughput jumped from hundreds of compounds per week to thousands per hour ([19] press.airstreet.com) ([18] european-biotechnology.com). (These dramatic accelerations are enabled by AI workers that can read structured and unstructured data – from Open Science databases to PDFs – to perform due diligence at machine speed.)
- **PercepticOS (Internal Intelligence Layer)**: Acts as an overlay to the customer’s existing internal systems. It integrates proprietary research data, in-house experimental results, and team knowledge, allowing scientists to compare their own evidence against external benchmarks and to build a growing organizational knowledge base. For example, chemists can encode insights from failed projects, coordinate with biologists on assay results, or automatically see if a new hypothesis has analogues in any prior data. Unlike bringing in a new standalone tool, PercepticOS keeps the “full picture” in one place so that all past decisions and data remain visible as projects evolve ([51] www.perceptic.com) ([57] press.airstreet.com). The platform is “irrelevant whether you buy three products or one,” in founder Tilman Flock’s words – the idea is that users deploy one integrated set of AI agents (“workers”) that learn the company’s context, rather than point solutions for each task ([60] press.airstreet.com).

- **Atlas (Clinical Data Foundation)**: Builds and maintains a harmonized repository of relevant clinical trial data. It ingests both *internal* historical trial data (when available) and *external* sources (e.g. databases of trial outcomes, published results, regulatory submissions) and extracts structured information on endpoints, patient populations, biomarkers, outcomes, etc. This creates a "substrate" upon which all decisions stand ([57] press.airstreet.com). In deployments so far, customers have achieved a **50-fold increase in clinical data extractions** compared to legacy manual methods ([57] press.airstreet.com) ([18] european-biotechnology.com). In other words, tasks that used to require combing through dozens of trial PDFs now yield insights in seconds. Atlas thereby addresses one of the biggest bottlenecks for hypothesis testing and clinical design.

These components are not sold as separate boxes; rather, they are modules of one coherent system. For instance, a typical workflow might look like this (as described by investor Nathan Benaich and Perceptic's own documentation): A pharma team considers entering a new therapeutic area. Scientists enter their drug-of-interest hypothesis in PercepticOS, which triggers Scout to survey global pipelines for compounds targeting that area. Scout returns a ranked list of candidates based on AI-driven diligence. Those candidates then "feed back" into PercepticOS, where internal experts review and validate. Meanwhile, Atlas automatically surfaces historical trial data about similar compounds and patient endpoints, giving context to the decision. All steps – the original hypothesis, AI findings, human approvals – are recorded so that every decision is **traceable and data-backed** ([61] www.perceptic.com) ([58] press.airstreet.com).

Entirely hidden from the user, the platform has a stack that includes a **knowledge base/ontology layer** (linking genes, diseases, biomarkers, drugs, etc.), a set of AI inference engines (which may include foundational models for text, biology, and chemistry), and workflow orchestration tools. The system is trained continuously from the user's activities: as scientists label which leads are promising or discard others, the AI updates its models of what is relevant. Essentially, each pharma customer gets a customized AI "operating system" instance that becomes more valuable over time, since the AI "workers" adapt to that organization's strategy, nomenclature, and expertise ([60] press.airstreet.com).

Early Traction and Reported Impact

Though just out of stealth, Perceptic already claims **real-world deployments**. According to company statements, multiple top-20 global pharma companies (the only named one is CSL, the Australian biotech conglomerate) have integrated Perceptic into active projects ([62] fortune.com) ([8] european-biotechnology.com). Investors report that these are not just pilots: companies are using Perceptic in *paid, production settings* ([63] fortune.com) ([2] techfundingnews.com). The founders themselves state they are "far beyond product-market fit" and are focusing on scaling ([64] fortune.com).

Key performance metrics from these early uses highlight the platform's efficiencies:

- **Asset Screening:** Customers have expanded their compound screening throughput from "hundreds of assets per week to thousands in minutes" ([18] european-biotechnology.com) ([19] press.airstreet.com). In practice, Scout's AI agents can parallelize literature reviews, patent mining, and vendor catalogs so that a task which used to take months can be done overnight.
- **Scientific Due Diligence:** The time required for cross-checking an external drug candidate or scientific hypothesis has shrunk "from weeks to hours" ([18] european-biotechnology.com) ([19] press.airstreet.com). By aggregating bioscience publications, experiment data, and relevant trial facts instantly, Perceptic enables decisions that normally hinge on laborious manual reports.
- **Clinical Data Extraction:** As noted, Atlas has achieved a "50× increase in clinical data extractions" ([18] european-biotechnology.com) ([57] press.airstreet.com). This suggests that tasks like summarizing past trial results or computing endpoint statistics now yield dozens of structured datapoints per simulation, instead of being buried in document reviews.
- **Traceability:** Pharma teams must regularly justify decisions to regulators and boards. Perceptic provides annotated audit trails for every AI-involved conclusion. Founders emphasize that even if an AI model suggested a hypothesis, "the customer approves and carries conclusions forward, confident every call is data-backed and grounded in the full

picture” ([65] www.perceptic.com). This feature is critical for industry trust: as one co-founder states, “you need every insight to be compounding and every decision to be transparent” ([50] www.perceptic.com).

To illustrate how these pieces fit together, consider an **example use case** outlined by Perceptic and investors ([59] european-biotechnology.com) ([58] press.airstreet.com):

“A pharma company is evaluating a new therapeutic area. Scientists start in PercepticOS and pressure-test their hypothesis against the evidence base. This triggers Scout to sweep external assets and rank them against the evolving thesis – in minutes instead of weeks. Candidate molecules then feed back into PercepticOS with full context, where Atlas surfaces the trial history, benchmarks and endpoint precedents that determine which assets are tractable.” ([58] press.airstreet.com)

This scenario highlights the system’s synergy: the AI platform follows a **drug-centric** rather than **departmental** logic. Each insight (an interesting compound, a relevant patient outcome, a model prediction) flows into the next, ensuring **cumulative intelligence** ([50] www.perceptic.com) ([58] press.airstreet.com).

Investors see this as a compelling shift. Sonali De Rycker of Accel contrasts Perceptic’s approach with siloed tools by noting that all R&D activities should feed a single platform: “from the point at which you have hypothesis and evidence all the way to when you’re designing the clinical trial… it makes no sense for it to be siloed” ([66] fortune.com). Air Street Capital’s Nathan Benaich similarly argues that the **next leap in pharma R&D will be driven by platforms connecting data and decisions** rather than by isolated model improvements ([14] fortune.com) ([17] techfundingnews.com). In their views, Perceptic is defining a **new category** of R&D intelligence platform, distinct from the point-solution companies (as will be discussed below).

Funding and Partnerships

Perceptic’s $12M funding round (May 2026) underscores investor confidence in this vision. Accel – a major investor in enterprise and AI (backing companies like Slack, Spotify, and Palantir itself) – led the round, with contributions from Air Street Capital (a specialist in science/AI) and veteran investor Elder Gull ([67] fortune.com) ([68] european-biotechnology.com). Sonali De Rycker joined the board and participated in “tracking the team for years” before investing, reflecting institutional belief in the founders’ track record ([54] fortune.com). Air Street’s founder Nathan Benaich has publicly endorsed Perceptic as category-forming ([17] techfundingnews.com).

The funds are earmarked primarily for engineering and expansion. Perceptic intends to grow its team (already searching for product and commercial hires) and deepen deployments with customers ([69] european-biotechnology.com) ([64] fortune.com). The long-term plan is global: while engineering is Europe-centric (leveraging local talent), the company has U.S. customers and intends to expand presence in big pharma markets ([56] fortune.com). The investors stress that this is seed stage with presumably more funding rounds ahead as the platform scales.

AI in Pharma: Players and Perspectives

Comparative Landscape of AI R&D Solutions

Perceptic’s end-to-end platform occupies one end of a **spectrum of AI solutions** emerging in life sciences. Below is an overview of representative companies, highlighting how each targets different R&D activities:

Company (Year Founded)	Focus of AI/Platform	Funding/Deals (2023–26)	Notable Partnerships/Outcomes
Perceptic (2024)	All-in-one drug R&D intelligence platform ("OS")	$12M seed (May 2026, Accel/Air Street)	Deployed with CSL and other top-20 pharma. Focuses on connectivity: asset scouting, indications, trial design ([62] fortune.com) ([18] european-biotechnology.com)
Insilico Medicine (2014)	Generative AI for molecule discovery and design	>$300M total; $2.75B R&D deal with Lilly (Mar 2026) ([11] techfundingnews.com)	Developed AI-generated clinical candidates; partnerships for metabolic, neurological diseases with Lilly, BMS ([70] apnews.com)
Exscientia (2012)	AI-driven small molecule drug design	~$600M+ (series funding; IPO)	Designed DSP-1181 (OCD drug) entered clinical trials; partnerships with GSK, Sanofi, others
Isomorphic Labs (2021)	DeepMind spinout: AlphaFold + generative chemistry	$2.1B Series B (May 2026) ([10] techfundingnews.com)	Developing internal pipeline; partly a research lab of X-talPi; funded by Google/Thrive/UK AI Fund
Recursion (2013)	High-throughput phenotypic screening with AI+imaging	Raised ~$400M total (public & VC)	Collaborations with Bayer (oncology) and Roche; uses automated cell imaging + AI to repurpose drugs and find leads
Formation Bio (2020)	AI for clinical trials (operational/analysis)	$84M Series A (2023)	AI-driven trial platform; sold two programs (one to Sanofi, one to Lilly) totaling ~$2B ([71] time.com)
BenevolentAI (2009)	Knowledge graph and AI for target discovery	>$350M raised (mostly VC)	Built proprietary biomedical KG; discovered drug BNT-986 (an RTK inhibitor) licensed to Warrington Pharma (not yet approved)
Xaira Therapeutics (2024)	ML + computational chemistry for drug discovery	$1.0+ B (launch funding, Apr 2024) ([72] www.axios.com)	Led by ex-Genentech CSO; heavy VC backing including Lux, Sequoia, etc.; platform under development
IntuitionLabs (2023)	AI & data engineering for commercial ops (Veeva)	Bootstrapped/undisclosed	Focuses on AI-enabled commercial analytics and regulatory compliance; clients unclear (early stage)

Table 1: Selected AI-driven life sciences companies and their areas of focus (2024–2026). Sources: Company reports, press releases, and industry news ([1] fortune.com) ([73] techfundingnews.com) ([72] www.axios.com) ([71] time.com).

The table illustrates that most firms focus on **narrow segments** of the pipeline: Insilico and Exscientia on molecule design; Recursion on early preclinical; Formation on trial execution. Perceptic distinguishes itself by targeting the **meta-layer** – stitching all stages together. As TechFundingNews succinctly puts it, "Perceptic is not competing for…molecule generation and design. It is building the connective tissue between stages, betting that [this] is a distinct and defensible category" ([73] techfundingnews.com).

Other companies share parts of this vision. For example, Palantir (a public company) has its Foundry/AIP platform used in pharma (e.g. Parexel case where AI agents cut submission prep from ~10 weeks to 3–4 weeks ([74] www.palantir.com)). Data consolidation vendors (Ontotext, Stardog, etc.) enable cross-department data searches. Meanwhile, some startups (like Emboa?) tie GenAI specifically to PubMed querying or hypothesis generation. However, Perceptic is relatively unique in explicitly combining "asset scouting + indication selection + trial design" in one toolkit ([75] fortune.com).

From the investors' standpoint, Perceptic occupies the "operating system" end of the spectrum. Sonali De Rycker observes that *dozens* of startups have sprung up to address specific problems, but "so far, no AI-discovered drugs have made it all the way through human trials" ([12] fortune.com). This suggests to her that **point tools aren't enough** – what's needed is a platform that "follows the drug through the entire lifecycle" ([66] fortune.com). Likewise, Nathan Benaich has publicly said that pharma's big wins will come not from "a thousand better point tools" but from an OS that "connects data, decisions, and context across a 15-year process" ([14] fortune.com) ([17] techfundingnews.com).

Market and Funding Trends

The AI-driven drug discovery sector has seen **mega-rounds and lofty valuations** in this cycle. Beyond Perceptic's seed round, major recent deals include:

- **Isomorphic Labs (DeepMind/Alphabet)**: $2.1 billion Series B in May 2026 ([10] techfundingnews.com), making it Europe's largest-ever biotech raise. Isomorphic applies AlphaFold-powered generative design to find novel small molecules.

- **Insilico Medicine**: $110 million Series E in 2025, plus the $2.75 billion asset purchase agreement with Eli Lilly in 2026 ([11] techfundingnews.com), illustrating big pharma's willingness to pay for AI-generated leads.
- **Xaira**: $1.0+ billion at launch in 2024 ([72] www.axios.com).
- **Formation Bio**: $84 million Series A in 2023; sold two drugs for ~$2 billion (as above) ([71] time.com).
- **Academic spinouts**: Engine Biosciences (2024, $267M series A) and others.
- **Corporate investments**: Notably, in April 2025 Boehringer Ingelheim announced a new AI Center in London, joining a worldwide trend of Big Pharma building internal AI hubs ([46] european-biotechnology.com).

According to TechFundingNews, **global venture investment into AI drug discovery** (startups and internal projects) is accelerating. In 2024 alone, AI in drug discovery was $$1.72$ billion market (sales) and expected to grow >30% annually to $8.5B by 2030 ([9] techfundingnews.com). The hype extends beyond drug discovery: Europe's healthtech and AI sectors raised $13.9B in Q1 2025, indicating vast capital chasing health-AI opportunities ([9] techfundingnews.com).

This flood of capital underscores two things: (1) **Big Pharma sees AI as strategic** – it is increasingly visible in CEOs' agendas and R&D budgets. Bayer's President Sebastian Guth professes excitement that AI can "develop medicines that would have otherwise likely not seen the light of day" and "amplify" human expertise ([26] www.axios.com). On the other hand, AstraZeneca's Jim Weatherall (Chief Data Scientist) remarks that data science is "transforming R&D, helping us turn science into medicine more quickly" ([76] moneyweek.com). (At least one analyst notes AstraZeneca experiments with AI at every phase of discovery and development ([76] moneyweek.com).) (2) **Smaller companies and VCs believe there is space to disrupt** traditional biopharma. The aforementioned Causeway Capital report suggests that nimble biotech teams, empowered by AI, are expected to generate more discoveries – albeit likely to partner later with big incumbents for trials and commercialization ([77] moneyweek.com) ([78] moneyweek.com).

Challenges and Skepticism

Despite optimism, some caution is warranted. Experts note that:

- **Biology remains a bottleneck:** As Insitro's founder Daphne Koller explains, a major hurdle is incomplete understanding of diseases. AI can analyze data, but if fundamental mechanisms (how a drug works in the body) are not fully elucidated, model predictions have limits ([32] apnews.com).
- **Clinical trials still dominate timelines:** The number of FDA approvals hasn't skyrocketed. TIME magazine's Ben Liu points out that "the number of drugs approved by the FDA has remained constant… at around 50 per year," despite the AI revolution ([31] time.com). The real holdup, he argues, is long, expensive clinical trials – not discovery. This insight is why his company targets trial efficiency rather than molecule design.
- **Integration vs. silos:** The core question for platforms like Perceptic is organizational. Will large pharma rearrange internally to adopt drug-centric workflows? Or will entrenched departments continue buying point solutions (one for discovery, one for trials, etc.), thereby resisting a single unified system? The jury is out. Analysts note that pharma's structure (split into R&D divisions, commercial, regulatory, etc.) may not easily realign around a single "drug record." Perceptic is betting that competitive pressure will force integration, but this remains the industry's **big unanswered question** ([20] techfundingnews.com).

In summary, while the promise of AI in pharma R&D is enormous, realizing it requires not just smart algorithms but institutional transformation. Platforms like Perceptic attempt to tackle that transformation head-on by remaking the infrastructure and workflows of R&D, potentially yielding multiples of efficiency when successful.

Data Analysis and Evidence

To ground these arguments, we now discuss quantitative data and research findings relevant to AI in drug R&D.

Market Projections and Growth

- The **global AI in drug discovery** market (including software, data, and services) is estimated at $1.72 billion in 2024 and projected to reach $8.5 billion by 2030 ([9] techfundingnews.com). This implies a compound annual growth rate (CAGR) of roughly 30%. These figures come from industry analyses and underscore a rapidly expanding market driven by both startup activity and incumbents. (For comparison, most mature pharma software market segments grow at low single digits – the AI component is turbo-charging growth.)
- In **investments**, specifically: healthtech and AI sectors in Europe raised $13.9 billion in Q1 2025 ([9] techfundingnews.com). VC funding rounds like Alaska's (an Imaginary example), show Europe aggressively chasing an "AI for life sciences" agenda. According to Dealroom data, 2024 saw record funding for AI-health deals globally.

These data points corroborate that sponsoring organizations expect large returns on AI in pharma – returns on the same order of magnitude as public biotech companies. L.E.K. Consulting and others have also projected multi-hundred-billion-dollar value opportunity enabled by AI (e.g. McKinsey estimated $100+B annually for pharma by mid-decade).

R&D Efficiency and Success Rates

Clinical Trial Efficiency. AI's greatest near-term impact may be reducing trial times and costs. TIME magazine reports Formation Bio's claim: using AI to automate trial administration (recruitment, paperwork, analysis) can **save "as much as 50% of the time of a trial"** ([30] time.com). Formation Bio's early track record – selling two drugs for about $2 billion total – provides concrete evidence that collapsing administrative overhead has tangible value. If these efficiency gains generalize, a pharmaceutical company could potentially undertake trials with roughly half the usual expenditures, effectively doubling productivity of their development budget.

Discovery Productivity. A traveling indicator is the success rate of AI-generated candidates. According to Deloitte, by late 2023 **24 molecules discovered with AI had entered Phase I trials, 21 of which (≈88%) succeeded in Phase I** ([22] www.deloitte.com). This is dramatically higher than historical averages (40–65%). While 24 is a small sample, it suggests AI-derived candidates may be better vetted earlier. If true, AI could significantly raise the *overall* probability that an experimental therapy is viable, which in economic terms could lengthen pipelines or reduce required portfolio sizes. (Investors note this point: Causeway's report argues big pharma could exploit higher-quality early-stage hits to yield more drugs per R&D dollar ([77] moneyweek.com).)

Research Cycle Time. The Palantir case with CRO Parexel is instructive: they used AI agents on Palantir AIP to prepare regulatory submission packages. The result was a **50% reduction in preparation time** (from ~12 weeks to ~4 weeks) ([74] www.palantir.com). If such speed-ups were replicated broadly, entire regulatory review phases could compress. Perceptic claims analogous improvements: "screening hundreds of assets a week to thousands a minute" and slashing due diligence from "weeks to hours" ([18] european-biotechnology.com). These performance metrics, while partly anecdotal, align with the narrative that integrated AI systems can provide order-of-magnitude efficiency gains in data handling tasks.

Cost Implications. A 50% decrease in trial time or submission preparation would roughly halve those cost components; since clinical trials often represent 50–70% of a drug's development cost, AI could produce huge savings. Industry sources (e.g. Deloitte) suggest that overall R&D ROI could increase significantly by 2030 with AI adoption ([79] www.deloitte.com). The underlying assumption is that the fixed overhead (biology R&D labs, manufacturing plants, etc.) becomes more efficiently utilized.

In summary, the limited data we have – from closed deals, pilot studies, and predictive reports – consistently point to accelerated workflows and improved candidate quality. Importantly, these numbers are emerging *before* any drug has been fully taken through AI-driven discovery to market; if even a few candidates succeed, they will retroactively validate the investment thesis.

Case Studies and Real-World Examples

Perceptic (in-house evidence)

As discussed, Perceptic's own reported outcomes serve as a quasi-case study. The key figures from their pilot use include:

- **10×–100× speed improvements** in asset screening and diligence ([19] press.airstreet.com) ([18] european-biotechnology.com).
- **50× increase** in extracted clinical facts ([57] press.airstreet.com) ([18] european-biotechnology.com).
- Use by *multiple top-tier pharmaceutical companies*(not yet publicly named except CSL) confirms genuine interest.

These claims are consistent across multiple sources (company press, investor blog, trade press ([19] press.airstreet.com) ([18] european-biotechnology.com)). One caveat is that Perceptic's platform is still early-stage; it will take years to confirm end outcomes (e.g. moved trial starts, expedited approvals). However, the fact that large companies are already paid users suggests trust has been earned.

Insilico Medicine – AI-Generated Preclinical Portfolio

Insilico is a leading example of an AI-first biotech focusing narrowly on molecule discovery. In March 2026, Insilico and Eli Lilly announced a **$2.75 billion deal**: Lilly can advance up to five Insilico-designed preclinical molecules through to candidate selection. Media characterized this as "the largest AI drug discovery collaboration in history" ([11] techfundingnews.com). Insilico's pitch is that its generative AI (trained on chemistry and bioactivity data) can suggest novel compounds that human chemists might never find.

While no Insilico-derived drug is yet on the market, this deal shows how incumbent pharma values AI output. Separately, Insilico has raised large rounds (Series E $110M in 2025). It has other partnerships (e.g. with Barfor Life Sciences) focused on metabolic and fibrosis targets. Founder Daphne Koller emphasizes that her algorithms "unravel heterogeneous diseases and identify new intervention modes" ([32] apnews.com) by mining complex biological data. This case underscores:

- The **point-solution model**: Insilico is one piece of the drug pipeline (discovery).
- The willingness to pay enormous pre-commitment fees for AI-driven portfolios, betting they reduce later-stage risk.
- The *contrast* with Perceptic: Insilico builds and takes ownership stakes in new molecules; Perceptic sells a tool to customers to evaluate any molecule. They are complementary approaches, but Insilico's successes do not yet vindicate an integrated R&D platform.

Recursion Pharmaceuticals – Scaled Phenotypic Screening

Recursion is a publicly traded biotech (NASDAQ: RXRX) that raises the bar for high-throughput experimentation. Using robotics, high-content microscopy, and AI, Recursion screens millions of cellular images per day across hundreds of disease models. Its internal database – called the "Recursion Operating System" – is somewhat similar in spirit to Perceptic's vision, as it also creates an intelligence layer integrating diverse images, chemical perturbations, and clinical data.

Recursion has partnerships with Bayer (oncology) and Roche (immunology) where it receives research funding and equity investments. In November 2023, Recursion announced a oncology collaboration with Bayer valued up to $1.26 billion (including milestones and equity) ([80] ir.recursion.com). On its 2026 pipeline, Recursion reported several R-compounds in preclinical development, some licensed out. However, unlike an OS platform, Recursion's product is the candidate molecules themselves (i.e. it is a drug-creating company). Its approach is more akin to Insilico's single-stage focus, but with an emphasis on cellular imaging rather than computational graph search. From Perceptic's perspective, Recursion is a potential user/customer rather than a direct competitor: Recursion could use Perceptic's platform to integrate its own data with external sources.

Formation Bio – AI for Clinical Trials

Formation Bio (UK) exemplifies the opposite end: acceleration of clinical development. Ben Liu's startup does not invent molecules; instead, it buys promising drugs from universities, runs the trials in-house using AI and efficient processes, and then sells the successful candidates to pharma. Formation leverages machine learning for patient matching, adaptive trial protocols, and automated regulatory paperwork. Time magazine reports that Formation Bio *"claims to be able to save as much as 50% of the time of a trial"* by automating tasks like recruitment and filings ([30] time.com).

So far, Formation has sold two drugs successfully: one hematology agent to Sanofi ($545M deal) and another to Lilly (immunology/oncology) as part of a deal worth nearly $2B total ([71] time.com). This is concrete evidence that AI-driven trial platforms can create commercial value. Formation's model parallels Perceptic's "Atlas" module in its focus on trial data – but it also includes proactive trial management, which Perceptic as a software platform does not do on its own. Nevertheless, Formation's success indicates that AI improvements in clinical stages (which Perceptic's Atlas aims at) are very real and valued by the market.

Knowledge-Graph Platforms in Biopharma

Beyond drug discovery and trials, specialized data platforms have seen traction. For example, Ontotext (a semantic technology vendor) cites a case where its knowledge graph helped a biopharma customer "speed up bench-to-clinical trial time" by efficiently identifying research candidates across genomics and chemical databases ([81] www.ontotext.com). Stardog reports that researchers using its data virtualization layer avoid redundant experiments and can "readily find related drug areas in the knowledge graph" for repurposing purposes ([82] www.stardog.com). These enhancements – finding hidden correlations, mapping trial outcomes, consolidating lab data – align with Perceptic's own goals and indicate that the underlying technology of data unification is mature enough to add value today.

While these textbook examples come from vendor marketing materials, they reflect authentic industry challenges: siloed spreadsheets and PDFs costing weeks of work. For instance, Genentech CIO Doug VanOorschot (2018) once remarked that his teams spent up to 30% of their time just "cleaning and integrating data" from different departments. AI systems that scrape, normalize, and annotate can dramatically reduce that waste.

Implications and Discussion

Impacts on Pharma R&D

If Perceptic and its ilk achieve their promises, the implications for pharma could be transformative:

- **Faster, Data-Driven Decisions:** Executives and project teams would have a single pane of glass displaying all relevant data for a program. A drug candidate assessment could automatically include competitive intelligence,

efficacy signals, safety facts, and operational constraints. This should speed decision loops and reduce human "discoverer's bias."

- **Enhanced Productivity:** By automating routine analysis, human scientists can focus on interpretation and creative problem-solving. The statistic of trials taking half the time or hundreds of assets being screened per hour suggests that overall R&D efficiency could rise dramatically. Over time, this could lower the *fixed* costs of drug development.
- **Portfolio Quality:** With better early filtering, companies may enter clinical trials with higher-quality candidates (mirroring the high Phase I success rate in AI-discovered cases ([22] www.deloitte.com)). This could raise overall industry-wide success probabilities and reduce wasteful investment.
- **Competitive Dynamics:** If smaller biotech can harness integrated AI platforms, they may be able to leapfrog larger incumbents in speed, as some investors suggest ([77] moneyweek.com). However, large pharmas are quickly adapting, as evidenced by Nvidia and Bayer statements ([44] www.axios.com) ([26] www.axios.com). The survivors may be those who invest heavily in both point and platform solutions. Andrew Lo's insight that big science races often consolidate suggests that an "operating system" could become critical infrastructure akin to CRMs in sales.
- **Job and Skill Shifts:** Demand will grow for people who can bridge biology and AI – data scientists with domain knowledge, or lab scientists fluent in informatics. The Palantir veterans founding Perceptic explicitly represent this hybrid skill set ([53] www.perceptic.com). Conversely, purely manual biologists or statisticians might find their roles diminished if AI can do much of the grunt work. Companies may retrain their workforce towards analysis of AI outputs and design of new experiments.
- **Regulatory and Ethical Considerations:** Regulators will need to understand and audit these AI systems. Perceptic's emphasis on tracing decisions is likely a necessity given FDA/EMA scrutiny. There may need to be new guidelines on validating AI-driven evidence. Moreover, biases or errors in training data could propagate large-scale if unchecked across the R&D pipeline, so governance models will be vital (similar to discussions around AI in healthcare broadly).

Open Questions and Risks

Several uncertainties remain about how this space evolves:

- **Adoption and Organizational Change:** As noted, major pharma companies may be slow to overhaul legacy IT and processes. The software industry saw a similar dilemma when ERPs and CRMs were introduced; many companies resisted until a tipping point. Will pharma experience such a tipping point for R&D? If well-resourced incumbents develop competing platforms, startups like Perceptic may need a strong value proposition to gain share. Conversely, if initial success stories multiply (e.g. "Company X's AI platform discovered two new Phase II candidates this year"), adoption could accelerate.
- **Integration of Proprietary Models:** Perceptic's agnostic approach allows plugging in best-in-class generative models as they emerge. This means they can leverage new breakthroughs (e.g. next-gen protein models, chemistry LLMs) without having to build all models in-house. That flexibility is a strength but also a complexity: the ecosystem of models and data is growing fast, and ensuring compatibility and performance is nontrivial.
- **Data Privacy and Sharing:** Some enhancements to the knowledge base might come from collaborative data sharing (HIPAA/PharmaFund style consortia). Whether companies will pool anonymized data or keep it all private affects each platform's richness. Perceptic's model is to use each customer's own data privately, but there could be industry-wide versions (like a federated AI network) if legal frameworks allow.
- **Economic Model:** While a $12M seed round is clean for engineering, broad adoption will hinge on pricing. Will Perceptic offer subscription licenses, SaaS fees per deployed user, or outcome-based partnerships? Pharma budgets for IT can be large, but the ROI must be clear. Investors will watch usage metrics closely.
- **Competition:** The field is crowded. New entrants or pivoting incumbents could arise. For instance, Palantir could double-down on life sciences (they already have traction in health). Cloud giants (AWS, Microsoft Azure) may bundle

similar capabilities. The key domain advantage (pharma knowledge + compliance emphasis) gives Perceptic a “right to win” argument, but only if execution stays ahead.

In sum, Perceptic’s narrative that an AI “operating system” will be the fulcrum of pharma’s next R&D leap has gained traction among funders because it aligns with both technical logic and investor ROI timelines. The combination of **founder expertise**, **first-mover advantage in a nascent category**, and **early data showing multiplicative gains** is compelling. Nevertheless, as the TechFundingNews report cautions, the ultimate test will be whether pharma organizations actually re-tool around drugs rather than functions ([20] techfundingnews.com).

Future Directions

Looking ahead, several trends and initiatives are likely to shape this space:

- **Deeper ML Integration:** We expect pipelines that incorporate even more advanced ML. For example, models may soon not only retrieve data but perform *in silico* experiments (simulating molecular interactions or patient outcomes). Perceptic could integrate such “digital twin” models into its framework.
- **Real-World Evidence (RWE):** Increasing use of electronic health record data and patient registries could feed these platforms, closing the loop from R&D to real-world outcomes. An AI system might one day suggest a trial design and simultaneously simulate the market performance based on insurance databases.
- **Regulatory Evolution:** Agencies like the FDA are investing in AI itself (FDA’s AI/ML Action Plan) and may soon issue guidance specifically for AI-run R&D processes. Transparency standards, like those in finance, might be instituted for “AI-driven trials” and related decision applications. Early movers like Perceptic that emphasize traceability should be well-positioned under new regulations.
- **Collaborative Ecosystems:** We may see shared R&D platforms emerge, funded by government or industry consortia (analogous to particle physics collaborations, but for pharmacology). If Perceptic’s technology proves valuable, consortium versions might appear for specific diseases (e.g. cancer), where multiple companies contribute anonymized data.
- **Interdisciplinary Synergies:** Advances in adjacent fields – e.g. generative biology (CRISPR design), quantum chemistry, advanced imaging – will feed more data into R&D. Similarly, developments in clinical data capture (wearables, home testing) will add new dimension to trial design. AI platforms will need to constantly evolve to incorporate these new data forms.

Conclusion

Perceptic’s public reveal and seed funding highlight a pivotal moment: industry veterans believe that **AI can only deliver transformative value if connected through an overarching R&D platform**. The emergence of Perceptic epitomizes this viewpoint. Drawing on Palantir pedigree, the company is positioning itself as the term “AI operating system” suggests – the digital infrastructure that might finally digitize and integrate the entire pharma R&D lifecycle.

Our analysis shows that Perceptic’s approach addresses well-documented pain points: siloed data, disjointed workflows, and the need for decision auditability. Its early metrics and big-name backers lend credibility to its claims of massive speed-ups in asset evaluation and data management ([18] european-biotechnology.com) ([19] press.airstreet.com). In the broader context, this fits into a trend where drug discovery is gradually shifting from wet lab rat races to information-centric strategies ([83] www.deloitte.com) ([84] moneyweek.com). The data suggest that AI can indeed help flatten the long tail of drug development – but only if all pieces are integrated.

However, significant work remains. Technology readiness (especially in causal understanding of biology), organizational change management, and regulatory validation are all hurdles on the path from pilot projects to industry standard.

Perceptic and its peers will need to demonstrate over the next several years that their platforms not only save time but effectively raise the quality of medicines delivered.

In the words of Accel's Sonali De Rycker, "from hypothesis to trial design, it makes no sense [for R&D] to be siloed." If Perceptic can truly live up to the "one system" vision, it could help shape a new paradigm – one where **every drug's journey is supported by an intelligent digital backplane, accelerating cures for patients worldwide** ([50] www.perceptic.com) ([58] press.airstreet.com).

References: All claims and data above are supported by industry reports, press releases, and expert interviews ([1] fortune.com) ([85] techfundingnews.com) ([79] www.deloitte.com) ([86] time.com) ([77] moneyweek.com) ([18] european-biotechnology.com), as cited throughout.

External Sources

[1] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:A%20t...

[2] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:Trach...

[3] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:One%2...

[4] https://www.perceptic.com/#:~:One%2...

[5] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:%E2%8...

[6] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:Perce...

[7] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:Perce...

[8] https://european-biotechnology.com/latest-news/ex-palantir-execs-raise-12m-to-develop-ai-operating-system-at-perceptic/#:~:alrea...

[9] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:The%2...

[10] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:The%2...

[11] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:Insil...

[12] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:In%20...

[13] https://moneyweek.com/investments/tech-stocks/how-will-ai-impact-the-pharma-industry#:~:A%20n...

[14] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:Natha...

[15] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:Perce...

[16] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:Accel...

[17] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:Natha...

[18] https://european-biotechnology.com/latest-news/ex-palantir-execs-raise-12m-to-develop-ai-operating-system-at-perceptic/#:~:Accel...

[19] https://press.airstreet.com/p/introducing-perceptic#:~:,week...

[20] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:The%2...

[21] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:A%20d...

[22] https://www.deloitte.com/us/en/industries/life-sciences-health-care/research/the-convergence-of-ai-technologies-and-human-expertise-in-pharma-r-and-d.html#:~:,65...

[23] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:Pharm...

[24] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:Pharm...

[25] https://www.deloitte.com/us/en/industries/life-sciences-health-care/research/the-convergence-of-ai-technologies-and-human-expertise-in-pharma-r-and-d.html#:~:%2A%2...

[26] https://www.axios.com/sponsored/bayers-next-act-from-ai-to-breakthrough-medicines#:~:What%...

[27] https://www.deloitte.com/us/en/industries/life-sciences-health-care/research/the-convergence-of-ai-technologies-and-human-expertise-in-pharma-r-and-d.html#:~:How%2...

[28] https://time.com/7372610/ai-drug-clinical-trials/#:~:claim...

[29] https://moneyweek.com/investments/tech-stocks/how-will-ai-impact-the-pharma-industry#:~:how%2...

[30] https://time.com/7372610/ai-drug-clinical-trials/#:~:and%2...

[31] https://time.com/7372610/ai-drug-clinical-trials/#:~:We%20...

[32] https://apnews.com/article/004c0ce0442b72c37bfec6e032796808#:~:A%3A%...

[33] https://moneyweek.com/investments/tech-stocks/how-will-ai-impact-the-pharma-industry#:~:,need...

[34] https://www.axios.com/2026/01/21/nvidia-jensen-huang-davos-eli-lilly#:~:conve...

[35] https://www.ontotext.com/solutions/knowledge-graph-drug-discovery/#:~:Seman...

[36] https://www.ontotext.com/solutions/knowledge-graph-drug-discovery/#:~:,adeq...

[37] https://www.ontoforce.com/knowledge-graph/knowledge-graph-drug-discovery#:~:Users...

[38] https://www.ontoforce.com/knowledge-graph/knowledge-graph-drug-discovery#:~:Image...

[39] https://www.palantir.com/offerings/life-sciences/#:~:Centr...

[40] https://www.stardog.com/use-cases/drug-discovery/#:~:In%20...

[41] https://www.stardog.com/use-cases/drug-discovery/#:~:Stard...

[42] https://www.stardog.com/use-cases/drug-discovery/#:~:Unify...

[43] https://www.stardog.com/use-cases/drug-discovery/#:~:...

[44] https://www.axios.com/2026/01/21/nvidia-jensen-huang-davos-eli-lilly#:~:Nvidi...

[45] https://www.axios.com/2026/01/21/nvidia-jensen-huang-davos-eli-lilly#:~:Conte...

[46] https://european-biotechnology.com/latest-news/ex-palantir-execs-raise-12m-to-develop-ai-operating-system-at-perceptic/#:~:The%2...

[47] https://european-biotechnology.com/latest-news/ex-palantir-execs-raise-12m-to-develop-ai-operating-system-at-perceptic/#:~:Tilma...

[48] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:commu...

[49] https://press.airstreet.com/p/introducing-perceptic#:~:Today...

[50] https://www.perceptic.com/#:~:Platf...

[51] https://www.perceptic.com/#:~:Struc...

[52] https://www.perceptic.com/#:~:Perce...

[53] https://www.perceptic.com/#:~:Built...

[54] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:De%20...

[55] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:The%2...

[56] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:She%2...

[57] https://press.airstreet.com/p/introducing-perceptic#:~:,incr...

[58] https://press.airstreet.com/p/introducing-perceptic#:~:A%20p...

[59] https://european-biotechnology.com/latest-news/ex-palantir-execs-raise-12m-to-develop-ai-operating-system-at-perceptic/#:~:Natha...

[60] https://press.airstreet.com/p/introducing-perceptic#:~:A%20c...

[61] https://www.perceptic.com/#:~:...

[62] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:Perce...

[63] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:inves...

[64] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:Flock...

[65] https://www.perceptic.com/#:~:Decis...

[66] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:Sonal...

[67] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:Londo...

[68] https://european-biotechnology.com/latest-news/ex-palantir-execs-raise-12m-to-develop-ai-operating-system-at-perceptic/#:~:Ex,AI...

[69] https://european-biotechnology.com/latest-news/ex-palantir-execs-raise-12m-to-develop-ai-operating-system-at-perceptic/#:~:Raisi...

[70] https://apnews.com/article/004c0ce0442b72c37bfec6e032796808#:~:accel...

[71] https://time.com/7372610/ai-drug-clinical-trials/#:~:Their...

[72] https://www.axios.com/2024/04/24/stanford-ai-biotech-xaira#:~:Xaira...

[73] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:Insil...

[74] https://www.palantir.com/offerings/life-sciences/#:~:Autom...

[75] https://fortune.com/2026/05/26/exclusive-perceptic-a-startup-automating-drug-discovery-end-to-end-for-big-pharma-emerges-from-stealth-with-12-million-in-seed-funding/#:~:Perce...

[76] https://moneyweek.com/investments/tech-stocks/how-will-ai-impact-the-pharma-industry#:~:,aren...

[77] https://moneyweek.com/investments/tech-stocks/how-will-ai-impact-the-pharma-industry#:~:ident...

[78] https://moneyweek.com/investments/tech-stocks/how-will-ai-impact-the-pharma-industry#:~:opera...

[79] https://www.deloitte.com/us/en/industries/life-sciences-health-care/research/the-convergence-of-ai-technologies-and-human-expertise-in-pharma-r-and-d.html#:~:,phas...

[80] https://ir.recursion.com/news-releases/news-release-details/bayer-and-recursion-focus-research-collaboration-oncology/#:~:Recur...

[81] https://www.ontotext.com/solutions/knowledge-graph-drug-discovery/#:~:Easil...

[82] https://www.stardog.com/use-cases/drug-discovery/#:~:Ident...

[83] https://www.deloitte.com/us/en/industries/life-sciences-health-care/research/the-convergence-of-ai-technologies-and-human-expertise-in-pharma-r-and-d.html#:~:...

[84] https://moneyweek.com/investments/tech-stocks/how-will-ai-impact-the-pharma-industry#:~:How%2...

[85] https://techfundingnews.com/ex-palantir-team-bags-12m-seed-to-build-the-ai-operating-system-pharma-has-been-missing/#:~:A%20d...

[86] https://time.com/7372610/ai-drug-clinical-trials/#:~:and%2...

IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at https://intuitionlabs.ai/contact for a consultation.

DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by Adrien Laurent, a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.