OpenDeSci | StGallen imt Business School

OPENDESCI RESEARCH INSIGHTS · READING DECK | OCTOBER 2026

When output is cheap, verification becomes scarce

The outlook for AI in science

How AI loops and graphs are reshaping discovery, and why attention, judgment and verification become strategic resources.

David Wang · Prof. Dr. Christian Abegglen · Dr. Florian Smeritschnig · Evan Luthra

In collaboration with StGallen imt Business School

opendesci.org

OpenDeSci | StGallen imt Business School

AT A GLANCE

AI is starting to run parts of the scientific method. Trust will decide who benefits.

01	**Evidence is accumulating**	Systems read, hypothesise, analyse, experiment and decide what to try next; proof points range from a Nobel Prize to a Phase III trial reported by Insilico Medicine.
02	**AI delivers where verification pathways are well specified**	Progress has been fastest where evaluation can be automated; lab and clinical verification often remains slow. When output is cheap, verification becomes scarce.
03	**Loops act; graphs connect and record**	The useful unit is no longer a model alone: it is a system with memory, evaluators and a record of provenance.
04	**Human capacities remain decisive**	Attention, originality, judgment, trust, responsibility and sovereignty shape what science pursues and accepts.
05	**Research organisations need a management response**	Roles, decision rights, evidence standards and escalation paths must change as AI agents take on research work.

Source: OpenDeSci Foundation, When Output Is Cheap, Verification Becomes Scarce: The Outlook for AI in Science (October 2026). The following pages number their sources, listed on pages 18 and 19.

OpenDeSci | StGallen imt Business School

01 | THE PROBLEM

Science produces more, but gets less per researcher, dollar and paper

USD 2.87 trn
Global R&D spending in 2024 (constant 2015 PPP), close to triple its level in 2000 [1]

>18×
More researchers now needed to keep Moore's Law going than in the early 1970s [2]

~80×
Fall in new drugs approved per inflation-adjusted billion dollars of R&D, 1950 to 2010 [3]

>70%
Of 1,576 researchers surveyed by Nature report failing to reproduce another scientist's experiment [4]



"AI can multiply the output of science. Whether it multiplies understanding is a management question."

Prof. Dr. Christian Abegglen
President, StGallen imt Business School

Science does not mainly lack intelligence. It lacks throughput, connection and verification.

Sources: WIPO Global Innovation Index (2025) [1]; Bloom et al., American Economic Review (2020) [2]; Scannell et al., Nature Reviews Drug Discovery (2012) [3]; Baker, Nature (2016) [4].

02 | THE EVIDENCE

AI delivers where verification pathways are well specified

Evidence is accumulating, but its strength varies: each card names its source type. Together, these milestones do not yet establish autonomous science.

PROTEIN STRUCTURE

200M

Protein structures predicted by AlphaFold2, honoured with the 2024 Nobel Prize in Chemistry [5]

Evidence: official prize announcement

MATERIALS

2.2M

Candidates predicted stable; 381,000 on the updated hull; 736 matched lab-made structures [6, 7]

Evidence: peer-reviewed; novelty contested

ALGORITHMS

48

Multiplications for 4×4 complex matrices: AlphaEvolve improved on Strassen's 1969 method [8]

Evidence: company disclosure

MATHEMATICS

35/42

Points: gold-medal standard at the 2025 International Mathematical Olympiad [9]

Evidence: company disclosure; IMO-certified

MEDICINE

Phase III

Insilico Medicine reported dosing the first Phase III patient for rentosertib, an investigational drug [11]

Evidence: company disclosure; Phase 2a peer-reviewed [10]

Progress has been fastest where evaluation can be automated. In materials science and medicine, verification remains experimental, staged and often slow.

Sources: Nobel Prize press release (2024) [5]; Merchant et al., Nature (2023) [6]; Cheetham and Seshadri, Chem. Mater. (2024) [7]; Google DeepMind (2025) [8, 9]; Xu et al., Nat. Med. (2025) [10]; Insilico (2026) [11].

A new industry is forming around automated discovery

Frontier labs, AI-scientist startups, drug developers and governments are building systems that close the loop from hypothesis to evidence.

FRONTIER LABS

Google DeepMind, OpenAI, Anthropic

- AI co-scientist: some hypotheses confirmed in early lab tests (Nature) [13]
- DeepMind: automated UK lab, announced Dec 2025 (company) [14]
- OpenAI: automated research intern (Sep 2026, self-reported) [15]

AI-SCIENTIST STARTUPS

Edison Scientific, Periodic Labs, Lila Sciences, Sakana AI

- Periodic Labs: USD 300M seed round, Sep 2025 (press) [18]
- Lila Sciences: USD 550M raised, valued above USD 1.3B (press) [19]
- Kosmos reads about 1,500 papers per run (preprint) [17]

AI DRUG DISCOVERY

Insilico Medicine and AI-native biotechs

- Rentosertib: Insilico reported dosing the first Phase III patient, Sep 2026 (company) [11]
- AI-found drugs to Dec 2023: 21 of 24 completed Phase I and 4 of 10 Phase II trials classed successful (small, non-exhaustive sample) [12]

GOVERNMENTS

United States, European Union

- Genesis Mission: executive order of 24 Nov 2025; DOE to mobilise its 17 labs (official) [20]
- EU RAISE pilot: EUR 107M, plus EUR 600M for compute (official) [21]

The gap: capital is flowing faster than shared infrastructure for verifying, connecting and explaining results.

Sources: Insilico [11]; Jayatunga et al. (2024) [12]; Gottweis et al., Nature (2026) [13]; company announcements [14–16]; Mitchener et al. (2025) [17]; TechCrunch [18]; Reuters [19]; US DOE [20]; European Commission [21].

From answers to loops: AI that acts, checks and tries again

AI systems create value when an evaluator closes the loop and decides what happens next.

REASONING LOOPS

Read, hypothesise, critique and refine

Closed by other agents, tournament-style ranking and human review

Example: Google's AI co-scientist [13]; Kosmos runs about 200 agent rollouts per objective [17]

EVALUATOR LOOPS

Propose code, algorithms or proofs

Closed by automated evaluators: tests, execution and formal checks

Example: AlphaEvolve improved on Strassen's 1969 algorithm [8]

PHYSICAL LOOPS

Design and run experiments with robots

Closed by measurements from the laboratory

Example: A-Lab made 36 of 57 target compounds in 17 days (count corrected in 2026) [24]

AI ACCELERATES SCIENCE FASTEST WHERE RELIABLE FEEDBACK CAN BE GENERATED

Mathematics and software — checked in seconds

Materials science — synthesis takes hours or days

Medicine — clinical loops close in years

METR: on software tasks, the length of tasks AI agents complete with 50% reliability has doubled about every 7 months since 2019 and about every 3 months since 2024 [22, 23]. Not a gauge of lab science.

Sources: Google DeepMind (2025) [8]; Gottweis et al., Nature (2026) [13]; Mitchener et al., arXiv:2511.02824 (2025) [17]; METR (2025, 2026) [22, 23]; Szymanski et al., Nature (2023, corrected 2026) [24].

Science is stored as documents, but it works as a graph

Three layers, not interchangeable. Provenance is a governance requirement, not a property of every knowledge graph.

GRAPH-BASED MODELS

90%

Of 1,380 targets: GraphCast, a graph neural network, beat the most accurate operational deterministic forecast system [25]. GNoME predicted 2.2 million candidate crystal structures [6].

KNOWLEDGE GRAPHS

17,080

Diseases in TxGNN's medical knowledge graph, 92% of them without an FDA-approved drug [26]. OpenAlex maps science itself as an open graph [27].

PROVENANCE GRAPHS

W3C PROV

A web standard since 2013 for recording who and what produced a piece of data [28]. In our assessment, the least developed layer in research, and the one that decides trust.

THREE QUESTIONS LOOPS CANNOT ANSWER ALONE

What is already known?

Without a map of prior work, a loop rediscovers the known and calls it new.

What contradicts what?

Failed replications, corrections and retractions must travel with every claim.

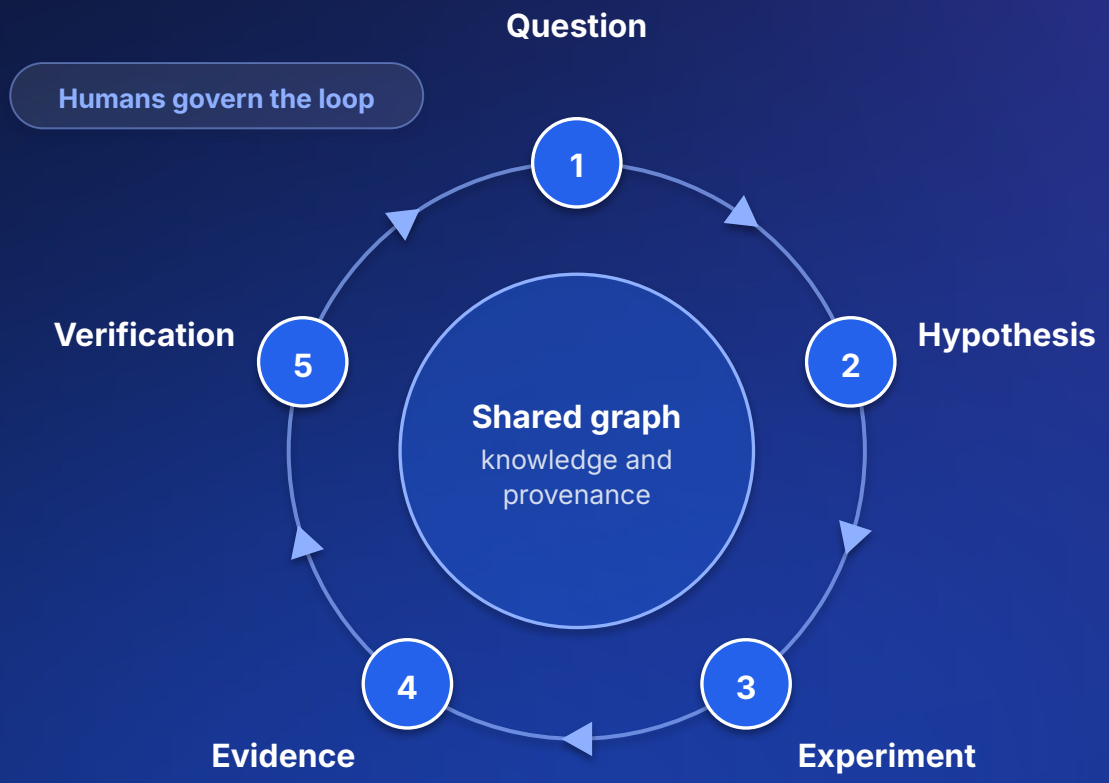
Where did this come from?

Every result needs a lineage: data, code, model and who signed off.

Sources: Merchant et al., Nature (2023) [6]; Lam et al., Science (2023) [25]; Huang et al., Nature Medicine (2024) [26]; OpenAlex documentation [27]; W3C PROV (2013) [28].

Loops act, graphs remember, humans govern

The next stage is a discovery engine linking models, evaluators, memory and governed workflows.



FIVE STAGES OF AI IN SCIENCE

The decisive step is from stage 2 to 3, when AI stops assisting scientists and starts running parts of research. It takes a new management system, not just a better model.



Sources: OpenDeSci Foundation analysis; five stages adapted from Wang and Abegglen (2026) [29] and Abegglen and Bleicher (2021) [30], a diagnostic framework, not a validated scale.

Five stages of AI in science: where do you stand?

Your stage is the number of the highest question you can answer with a yes.

STAGE 1
AI as search

YOU ARE HERE IF

Researchers use AI mainly to search, summarise and translate

STAGE 2
AI as assistant

YOU ARE HERE IF

AI drafts code, analyses or text; a person directs and checks every step

STAGE 3
In-silico scientist

YOU ARE HERE IF

An AI system generates and tests hypotheses on its own; people review the output

STAGE 4
AI as lab partner

YOU ARE HERE IF

AI decides which experiments to run next, and robots run them

STAGE 5
AI-native discovery

YOU ARE HERE IF

Several loops share one provenance-tracked knowledge graph and improve each other

84% of researchers use AI tools; 62% for research and publication tasks (Wiley, 2025; self-reported) [31]

47% vs 20% of respondents report scaling chatbots vs AI agents across the enterprise (McKinsey, 2026; self-reported) [32]

FIVE QUESTIONS TO PLACE YOURSELF

1 Do your researchers use AI mainly to search, summarise and translate?
2 Does AI write code, analyse data or draft text, with a person directing each step?
3 Does any AI system generate and test hypotheses on its own, with people reviewing?
4 Does AI choose the next experiments, and do robots run them?
5 Do several AI loops share one provenance-tracked knowledge graph?

Available surveys are consistent with widespread use at stages 1 and 2, but they do not directly measure this framework.

Sources: Wiley, ExplanAItions 2025 (October 2025, 2,430 researchers) [31]; McKinsey, The State of AI in 2026 (August 2026) [32]; authors' framework (page 8), not a validated classification.

OpenDeSci | StGallen imt Business School

08 | THE BOTTLENECK

When output is cheap, verification becomes the scarce resource

AI accelerates useful discovery, but it also accelerates error, fabrication and a narrowing of science.

BEYOND VERIFICATION

Attention, originality, judgment, trust, responsibility and data sovereignty remain scarce and require human and institutional governance.

100

Hallucinated citations in 50+ accepted NeurIPS 2025 papers, reported by the AI-detection vendor GPTZero [33]

21%

Of 75,800 ICLR 2026 peer reviews flagged as fully AI-generated by a detector [34]

−4.63%

Volume of topics studied as AI use spreads; 22% less engagement between scientists (observational) [35]

REBUILD TRUST AS INFRASTRUCTURE

01

Provenance by default

Every claim carries its data, code, model version and accountable humans.

02

Verification at machine speed

Automated checks of citations, statistics and code, plus rewards for replication.

03

Diversity by design

Funding and evaluation that reward data-poor questions, not just more output.

04

Transparency for the public

Sources shown, uncertainty stated, corrections visible.

Sources: GPTZero (January 2026) [33]; Nature news on the Pangram Labs analysis (November 2025) [34]; Hao et al., Nature (2026) [35].

09 | INTEGRATED MANAGEMENT

AI in science is a management-system question

Agentic research changes purpose, positioning and operating controls, not only the technical toolset.

NORMATIVE | ORIENTATION

Why use AI, and within which limits? Who remains accountable for each published claim?

WHAT MATURITY LOOKS LIKE

Humans stay accountable for every published claim. AI use is disclosed. Openness and provenance are institutional values.

STRATEGIC | POSITIONING

Where will we build or partner: data, evaluators, labs, graphs or talent?

WHAT MATURITY LOOKS LIKE

An explicit value thesis for AI in research, and a protected, diverse research portfolio.

OPERATIONAL | EXECUTION

Where do verification gates, disclosure, escalation and human sign-off sit?

WHAT MATURITY LOOKS LIKE

Provenance logged in every workflow, verification gates before publication, success measured by validated results.

Governance cuts across all three: who may let an agent run, on which budget, against which standard of proof.

Agentonomics lens [37]: explicit roles, decision rights, evidence requirements and escalation paths for people and agents. Without them, pilots are likely to multiply while knowledge production stays the same, a pattern Abegglen calls organised standstill [36].

Sources: Abegglen and Bleicher, Das Konzept Integriertes Management, 10th edn (2021) [30]; Abegglen (2021) [36]; Abegglen et al., TUM-I26129 (2026) [37]; used here as an analytical lens.

Open science first, DeSci where it beats the alternative

Blockchain may add value where it solves a real coordination or auditability problem. Conventional open infrastructure remains the default.

WHERE BLOCKCHAIN MAY ADD VALUE: ONE OPTION, NOT A NECESSITY

01 Provenance anchoring

Tamper-evident fingerprints of data, code, models and claims

02 Credit for verification

Portable records of who reviewed, replicated or corrected

03 Transparent funding

Grant and replication-bounty flows anyone can follow

04 Programmable payouts

Regulated stablecoin payments, where they create measurable value

> *The test: decentralise a component only where it solves a provenance, incentive or governance problem better than repositories, DOIs, signatures or open databases.*

Sources: OpenDeSci Foundation (2026) [38]; authors' analysis. Token market data: page 13.

OPENDESCI: THE PUBLIC INTERFACE

OpenDeSci aims to make science accessible to the next generation, with every explanation linked to its sources and its uncertainty stated.

> ***"AI will not replace scientists, and DeSci will not replace universities. What both can replace is gatekeeping that no longer scales: closed review, opaque credit and slow verification."***



Dr. Florian Smeritschnig
Founder, StrategyCase.com · Advisor, OpenDeSci

DeSci after the hype: from token prices to real-world value

The DeSci token market on 2 October 2026, what went wrong at Bio Protocol, and the view of Evan Luthra, Web3 and blockchain investor.

DESCI MARKET SNAPSHOT: 2 OCTOBER 2026

USD 0.3B

combined value of 58 DeSci tokens, under a third of BIO's value at its peak; five selected tokens shown [40]

OriginTrail	USD 175M
Bio Protocol	USD 78M
ResearchHub	USD 9M
VitaDAO	USD 3.5M
ValleyDAO	USD 0.7M
53 other tokens	USD 39M

CASE: BIO PROTOCOL (BIO)

~USD 1.2B ▸ **USD 78M**

estimated circulating market value at peak, 3 Jan 2025 [39, 41]

2 Oct 2026; price 96.6% below its peak [39]

WHAT WENT WRONG (OUR ANALYSIS)

- Priced on a listing and a narrative, not on results
- Low float at launch (about 39%), then nearly two years of unlocks: about 75% now circulates [39, 41]
- Little public evidence of users, revenue or results

WHAT IS NEW

A launchpad for AI research agents (BioAgents) [39]



Evan Luthra

Co-founder and General Partner, KOL Capital · Advisor, OpenDeSci

"DeSci's first wave priced the dream, not the science. … Decentralisation matters when it makes science more open and more trusted. The real asset is the innovation, not the token."

WHAT MATTERS NEXT **Real-world adoption · Real-world use cases · Decentralisation where it adds trust · Innovation in science**

Sources: CoinMarketCap, Bio Protocol [39] and DeSci category [40] (accessed 2 October 2026); 1inch token profile (2026) [41]. Peak value = all-time-high price × tokens circulating at launch.



David Wang at the AI Agents Summit 2026, Berlin

OpenDeSci | StGallen imt Business School

12 | OUTLOOK 2030

Faster where evaluation is automated, slower where it is not

FIVE FORECASTS, EACH WITH A DEADLINE

01 End-2028: at least 3 major funders or journals require machine-readable provenance

02 End-2030: 10,000+ new materials from autonomous labs, independently validated

03 End-2030: a regulatory decision on rentosertib; we expect approval [42]

04 End-2028: at least 5 of the 10 largest publishers check every submission automatically

05 End-2030: most US and EU adults report using AI monthly for science or health

> *"The opportunity is not to publish faster. It is to make knowledge verifiable, connected and understandable for everyone, from the scientist in the lab to the student on a phone."*
>
> **David Wang, Founder, OpenDeSci**

David Wang · Prof. Dr. Christian Abegglen · Dr. Florian Smeritschnig · Evan Luthra
In collaboration with StGallen imt Business School

Sources: authors' forecasts; rentosertib timeline from Bio-IT World (July 2026) [42].

From chatbots to agents to self-directed loops

The authors' view of where AI goes next.

WE ARE HERE

ERA 1 | 2022–2024

Chatbots

1.2B

weekly ChatGPT users, as reported by OpenAI in Sep 2026, after launch in Nov 2022 [43]

ERA 2 | 2025–2026

Agents

40%

of large firms report scaling agents in at least one function, up from 27%; smaller firms 22% [32]

ERA 3 | EMERGING, TO 2030

Self-directed loops on graphs

1 month

a human work-month of software tasks, done autonomously by 2030 if METR's trend holds [22]

ERA 4 | VISION, AFTER 2030

AI-native organisations

Level 5

OpenAI's reported top level: AI that can do the work of an organisation [44]

PEOPLE MOVE FROM ASKING TO GOVERNING

Ask and check → **Delegate and review** → **Set goals and standards of proof** → **Govern purpose and accountability**



"AI is growing more self-directed and making its own decisions. Through loops and graphs, it may soon overrule ours. Who sets its rules?"

David Wang, Founder, OpenDeSci

Limited and unreliable introspective access: in Anthropic's concept-injection experiments, Claude Opus 4.1 correctly detected and identified injected concepts in about 20% of trials under the best-tested conditions [45]. The study does not establish machine self-awareness.

Gartner: agentic AI in 33% of enterprise apps by 2028, yet over 40% of agentic AI projects cancelled by end of 2027 [46].

Sources: METR (2025) [22]; McKinsey, The State of AI in 2026 [32]; OpenAI DevDay (Sep 2026) [43]; Axios on OpenAI's five levels (Jul 2024) [44]; Anthropic introspection study (Oct 2025) [45]; Gartner (2025) [46].

14 | WHAT TO DO NOW

What leaders should decide now

ONE SEQUENCE FOR EVERY ORGANISATION

01
Choose one verifiable loop
Start where a credible evaluator already exists.

02
Define the standard of proof
State what evidence closes the loop and what remains uncertain.

03
Log provenance from day one
Record data, code, model versions, corrections and sign-off.

04
Assign human decision rights
Set budgets, stop rules, escalation paths and final accountability.

FIRST MOVES BY GROUP

RESEARCH LEADERS AND UNIVERSITIES

- Assess where you stand on the five stages
- Name an accountable owner for AI in research

FUNDERS AND POLICYMAKERS

- Fund replication, open data and open graphs as infrastructure
- Attach provenance rules to public AI-for-science programmes

COMPANIES AND INVESTORS

- Back loops with real evaluators and proprietary data, not demos
- Judge claims on prospective human evidence, such as controlled Phase II results; seek independent replication where feasible

SCIENTISTS AND COMMUNICATORS

- Use AI agents, but sign only what you can verify
- Explain findings with their sources and their uncertainty

The real question is not how much science AI can produce. It is how much of it we can trust, connect and understand.

Source: authors' synthesis of the article, drawing on the integrated management lens [30] and Agentonomics [37] (page 11).

OpenDeSci | StGallen imt Business School

ABOUT THE AUTHORS

The authors



David Wang

Founder, OpenDeSci · Managing Director, Wang Advisory GmbH · Senior Lecturer, StGallen imt Business School

Entrepreneur and management consultant with more than 10 years of international consulting experience at Deloitte and on projects for Bain & Company and Kearney. Has advised private equity funds and Fortune 500 companies on value creation.



Prof. Dr. Christian Abegglen

President, StGallen imt Business School · Advisor, OpenDeSci · Co-author of Agentonomics

Co-author and editor of Das Konzept Integriertes Management, the standard work on the St. Gallen concept founded by Knut Bleicher. Author of Unternehmen neu erfinden, a workbook against organised standstill.



Dr. Florian Smeritschnig

Founder, StrategyCase.com · Advisor, OpenDeSci

PhD, University of Vienna; Master's, Hong Kong University of Science and Technology. Previously McKinsey and Bitpanda. He is also a pilot.



Evan Luthra

Co-founder and General Partner, KOL Capital · Advisor, OpenDeSci

Web3 and blockchain investor. Angel investor since 2014 and early Bitcoin advocate; invests across Web3, AI and frontier technology. Forbes 30 Under 30.

In collaboration with StGallen imt Business School | opendesci.org

OpenDeSci | StGallen imt Business School

SOURCES

References

Numbered in order of first citation, page by page. All sources accessed between 30 September and 26 October 2026. Underlined links open the source.

1 WIPO Global Innovation Index team. Despite the odds, global R&D spending grew again in 2024. *WIPO* (2025). wipo.int

2 Bloom, N., Jones, C. I., Van Reenen, J. & Webb, M. Are ideas getting harder to find? *Am. Econ. Rev.* **110**, 1104–1144 (2020). doi:10.1257/aer.20180338

3 Scannell, J. W., Blanckley, A., Boldon, H. & Warrington, B. Diagnosing the decline in pharmaceutical R&D efficiency. *Nat. Rev. Drug Discov.* **11**, 191–200 (2012). doi:10.1038/nrd3681

4 Baker, M. 1,500 scientists lift the lid on reproducibility. *Nature* **533**, 452–454 (2016). doi:10.1038/533452a

5 Royal Swedish Academy of Sciences. The Nobel Prize in Chemistry 2024. Press release (9 October 2024). nobelprize.org

6 Merchant, A. et al. Scaling deep learning for materials discovery. *Nature* **624**, 80–85 (2023). doi:10.1038/s41586-023-06735-9

7 Cheetham, A. K. & Seshadri, R. Artificial intelligence driving materials discovery? *Chem. Mater.* **36**, 3490–3495 (2024). doi:10.1021/acs.chemmater.4c00643

8 AlphaEvolve team. AlphaEvolve: a Gemini-powered coding agent for designing advanced algorithms. Google DeepMind (14 May 2025). deepmind.google

9 Google DeepMind. Gemini Deep Think achieves gold-medal standard at the International Mathematical Olympiad (21 July 2025). deepmind.google

10 Xu, Z. et al. A generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis: a randomized phase 2a trial. *Nat. Med.* **31**, 2602–2610 (2025). doi:10.1038/s41591-025-03743-2

11 Insilico Medicine. First patient dosed in GENESIS-IPF-3, the Phase III trial of rentosertib (10 September 2026); ClinicalTrials.gov NCT07687459. insilico.com

12 Jayatunga, M. K. P., Ayers, M., Bruens, L., Jayanth, D. & Meier, C. How successful are AI-discovered drugs in clinical trials? *Drug Discov. Today* **29**, 104009 (2024). doi:10.1016/j.drudis.2024.104009

13 Gottweis, J. et al. Accelerating scientific discovery with Co-Scientist. *Nature* **655**, 487–496 (2026). doi:10.1038/s41586-026-10644-y; preprint arXiv:2502.18864 (2025).

14 Hassabis, D. et al. Strengthening our partnership with the UK government. Google DeepMind (10 December 2025). deepmind.google

15 OpenAI. Research acceleration: the view inside OpenAI (6 September 2026). openai.com

16 Anthropic. Claude for Life Sciences (20 October 2025). anthropic.com

17 Mitchener, L. et al. Kosmos: an AI scientist for autonomous discovery. arXiv:2511.02824 (2025).

18 TechCrunch. Former OpenAI and DeepMind researchers raise whopping $300M seed to automate science (30 September 2025). techcrunch.com

19 Reuters. AI lab Lila Sciences tops $1.3 billion valuation with new Nvidia backing (14 October 2025). Syndicated copy: finance.yahoo.com

20 U.S. Department of Energy. Energy Department launches Genesis Mission (24 November 2025); Executive Order 14363 (2025). pnnl.gov

21 European Commission. European Strategy for Artificial Intelligence in Science and RAISE (8 October 2025). ec.europa.eu

22 METR. Measuring AI ability to complete long software tasks (19 March 2025). metr.org

SOURCES

References

Continued from page 18. DOIs open the publisher's page via doi.org.

23 METR. Time Horizon 1.1 (29 January 2026). metr.org

24 Szymanski, N. et al. An autonomous laboratory for the accelerated synthesis of inorganic materials. *Nature* **624**, 86–91 (2023). doi:10.1038/s41586-023-06734-w; Author Correction, *Nature* **650**, E1 (2026).

25 Lam, R. et al. Learning skillful medium-range global weather forecasting. *Science* **382**, 1416–1421 (2023). doi:10.1126/science.adi2336

26 Huang, K. et al. A foundation model for clinician-centered drug repurposing. *Nat. Med.* **30**, 3601–3613 (2024). doi:10.1038/s41591-024-03233-x

27 OpenAlex. Developer documentation (2026). developers.openalex.org

28 W3C. PROV-Overview: an overview of the PROV family of documents. Working Group Note (30 April 2013). w3.org

29 Wang, D. & Abegglen, C. Integrated management in the AI era. White paper, Wang Advisory GmbH with StGallen imt Business School (26 August 2026). wangadvisory.com

30 Abegglen, C. and Bleicher, K. (2021). *Das Konzept Integriertes Management: Visionen – Missionen – Programme*. 10th fully updated and expanded edition. Frankfurt/New York: Campus. ISBN 978-3-593-51171-9.

31 Wiley. ExplanAItions 2025: The Evolution of AI in Research (7 October 2025). newsroom.wiley.com

32 McKinsey & Company. The state of AI in 2026: on the road to ROI (25 August 2026). mckinsey.com

33 GPTZero. GPTZero finds 100 new hallucinations in NeurIPS 2025 accepted papers (21 January 2026). gptzero.me

34 Naddaf, M. Major AI conference flooded with peer reviews written fully by AI. *Nature* **648**, 256–257 (2025). doi:10.1038/d41586-025-03506-6; Pangram Labs (18 November 2025). pangram.com

35 Hao, Q., Xu, F., Li, Y. & Evans, J. Artificial intelligence tools expand scientists' impact but contract science's focus. *Nature* **649**, 1237–1243 (2026). doi:10.1038/s41586-025-09922-y

36 Abegglen, C. *Unternehmen neu erfinden*, 2nd edn (Frankfurter Allgemeine Buch, 2021). ISBN 978-3-96251-087-9.

37 Abegglen, C., Knoll, A. C., Suckfüll, H. C. & von Grone, J. Agentonomics: Management im Zeitalter intelligenter Agentensysteme. TUM-I26129, TU München (2026). gimsg.ch

38 OpenDeSci Foundation. Website (2026). opendesci.org

39 CoinMarketCap. Bio Protocol (BIO): price, market capitalisation and supply. Accessed 2 October 2026. coinmarketcap.com

40 CoinMarketCap. Top DeSci tokens by market capitalisation. Accessed 2 October 2026. coinmarketcap.com

41 1inch. Bio Protocol (BIO) token profile (2026). 1inch.com

42 Proffitt, A. Insilico Medicine launches Phase III trial for AI-developed drug. *Bio-IT World* (8 July 2026). bio-itworld.com

43 OpenAI. DevDay 2026 recap (29 September 2026). openai.com

44 Axios. OpenAI nears "reasoning"-capable AI (15 July 2024). axios.com

45 Anthropic. Signs of introspection in large language models (29 October 2025). anthropic.com

46 Gartner. Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027. Press release (25 June 2025). gartner.com

IMPORTANT INFORMATION

Disclaimer

This publication is for general information only. It does not constitute investment, financial, legal or tax advice, or an offer or solicitation to buy or sell any security or token. Views are the authors' own and do not necessarily reflect those of their affiliated institutions. Forward-looking statements are not guarantees.

The authors founded, lead or advise OpenDeSci. They may hold financial interests in OpenDeSci and in companies, projects or tokens mentioned in this publication.

Version of 26 October 2026. Trademarks belong to their respective owners.
© 2026 OpenDeSci Foundation.