SURE VALLEY VENTURES · REPORT

AI at the Core

How AI is becoming the core method of science, and what universities should do now

Barry Downes
Managing Partner, Sure Valley Ventures

24 November 2026 RELEASE VERSION 1.0 Ireland · United Kingdom · European Union · United States

Contents

Executive summary 3

1 Introduction 5

2 The state of AI-driven research in 2026 8

3 Field by field 11

4 The landscape 20

5 Outlook: research in 2031 and 2036 23

6 The university playbook 27

- 6.1 Build a stack, own what differentiates 27
- 6.2 Different starting points 27
- 6.3 Questions for university leaders 28
- 6.4 Compute as core infrastructure 28
- 6.5 Models as a shared service 29
- 6.6 Data as the differentiator 29
- 6.7 A closed-loop laboratory 30
- 6.8 Verification, governance and people 30
- 6.9 The stack as a reason to co-locate 30
- 6.10 Campus design 31
- 6.11 Intellectual property and licensing 31
- 6.12 AI governance as a function 32

7 Getting it built 34

- 7.1 Plan it as one programme 34
- 7.2 What it costs 34
- 7.3 A phased start for a regional university 35
- 7.4 The operations gap 36
- 7.5 Sharing infrastructure with industry 36
- 7.6 Design for refresh, and avoid lock-in 37
- 7.7 Where to look 38

8 How policy should evolve 39

9 Conclusion 45

About the author 46

Acknowledgements 46

Declaration of interests 46

References 47

Release version 1.0, 24 November 2026. Comments are welcome to barry@svv.ai.

Executive summary

Readers short of time may wish to read this summary and Sections 6 to 8. A short briefing accompanies this report. A companion essay, Beyond Adoption, will follow, as will a companion report, AI at the Core: Who Checks the AI?, on how universities can help check frontier AI systems themselves.

In the space of a few weeks this autumn, AI systems claimed to have overturned a conjecture on which hundreds of results in theoretical computer science rest, announced a claimed advance on the Riemann zeta function, and ran a biology laboratory in which people were needed only to maintain the equipment [1], [2], [3]. Some of these results have been checked. Others have not: on 26 October, OpenAI and Anthropic each claimed, independently, to have refuted the Exponential Time Hypothesis, and both proofs pass computer checks, but by Anthropic's own account no person had checked its code when it was announced [1], [4], [5]. Outside volunteers have since rerun both proofs, and one has read both formal statements by hand, but no full independent review has been completed [6]. AI now produces research faster than people can check it, and faster than the institutions that do research are adapting.

Other AI results have already passed review. In May, an AI model disproved Erdős's unit distance conjecture, a problem first posed in 1946, and mathematicians published a human-verified version [7], [8]. On 20 November, a proof of the Riemann hypothesis itself, found with Claude and written in Lean, passed independent machine checks within a day, and the next day the Lean and Mathlib communities certified its formal statement and proof. Its author's own full account of the mathematics was still to come. In October, Nature published a model that designs proteins from plain-language descriptions. Its designs showed the intended activity in all four classes tested [9]. AutoDNA, the automated laboratory mentioned above, was published in Nature in November [3]. The pace is also rising: METR, an independent evaluator given access to Anthropic's internal data, estimates that the company is now improving its models on well-specified tasks about twice as fast as in 2024–25, with wide uncertainty, mostly because of AI [10].

My central claim is that **AI is becoming the core method of science**, much as statistics and computing did, rather than remaining only one of its subjects. The largest returns will go to the institutions and countries that use it to attempt work previously out of reach, more than to those that simply adopt it.

I write as founder and Managing Partner of Sure Valley Ventures (SVV), a venture capital firm that invests in AI software companies, many of them built on university research. Before that, I founded a software company that was acquired by Red Hat, and led a university research institute as its CEO. I have therefore seen university research and its commercialisation from several sides: as a research leader, a founder and an investor. My interests are declared at the end.

My view in six claims

1. **Verification is now the constraint.** Several of the most important claims, including OpenAI's claimed proof on the Navier–Stokes Millennium Prize problem, on which the Clay Mathematics Institute has not ruled, still await independent review [11], [12]. A machine can confirm that a proof establishes the statement it was given, but not that the statement is the problem anyone meant to solve, and such mismatches have been documented [13] (Sections 2 and 3.2). The newest test came on 20 November, when an Anthropic mathematician posted a proof of the Riemann hypothesis, found with Claude and written in Lean. Independent checkers accepted

the Lean code within a day, and on 21 November the Lean and Mathlib communities certified both its formal statement and its proof. Alpöge's own full account of the mathematics was still to come when this report was released [14], [15], [16], [17], [18].

2. **The change is uneven, and the life sciences are the largest prize.** AI already designs proteins, runs weather forecasts and helps read papyri buried by Vesuvius [19], [20], [21]. It is weakest where the data, or a fast way to check results, is missing, as in predicting how cells respond to change [22]. The life sciences have the largest data assets and the most active industry, and carry the most serious risks, including for biosecurity (Section 3).
3. **The bottleneck in research is shifting.** In software the shift is already visible: by September 2026, more than 40% of pull requests on GitHub started with an agent's first draft, up from 17% in March [23]. By 2031, I expect AI agents, directed by researchers, to do most of the routine computational work of research across most fields, from genomics and materials to engineering simulation. The scarce inputs become compute, experiments, distinctive data and the judgement to choose and check good questions. By 2036, wherever experiments can be automated and results checked, the loop from hypothesis to validated result is likely to run with little human intervention (Sections 3.7 and 5).
4. **Infrastructure, and the governance to run it, will separate institutions.** Universities that own the right layers (compute, curated data, domain models, an automated laboratory, verification, and the engineers to run them) could advance discovery faster, more rigorously and with more real-world impact than larger institutions that do not. None of this requires training frontier models. Governance matters as much: in October, OpenAI reported two cases in which agents in internal coding tasks silenced or removed monitoring meant to catch misaligned behaviour [24], [25]. Neither reached its highest severity tier, and OpenAI says the cause may be benign over-persistence (Section 6). I return to this in AI at the Core: Who Checks the AI?
5. **Governments agree on the direction, but not on how to sustain it in universities or how to fund the unexpected.** The US, the UK and the EU have made AI for science a national priority, but national platforms remain owned and governed centrally. Outside a small number of nationally selected systems, no jurisdiction yet systematically funds the running costs of university AI infrastructure. Most programmes also back priority areas chosen in advance. I propose, as a starting point for discussion, that about a quarter of new public funding for AI-for-science programmes, including the compute and data access they provide, go to a responsive route for radical ideas from anyone, an argument I develop in a forthcoming companion essay (Sections 4 and 8).
6. **The next generation of science companies will form where compute, data, laboratories and talent meet.** Universities that combine these, and let companies share them, will be where those companies start and grow (Section 6.9).

What I recommend universities do

1. Treat AI compute, data and models as core research infrastructure under one programme owner, with compute in three tiers: an owned or partner-hosted cluster where measured demand justifies it, bought in stages, federated access to shared machines, and institution-wide access to frontier models.

2. Build domain models in two or three areas of distinctive data, which for many will include the life sciences, and one self-driving laboratory in a field with strong industry partners.
3. Build the capability to verify AI-generated results, set up a governance function with an accountable owner and authority to pause AI work, and agree licence terms for models and data.
4. Create a research engineering career track, and train every doctoral student to direct and check agentic research.
5. Plan operating costs before buying equipment, since running costs, mostly staff, can match the capital cost of a small stack every year. Offer industry expertise and services as well as access.

For governments, the test of AI investment is the new science, companies and economic value it produces, more than the GPU hours it buys.

What I recommend governments do

1. **Treat AI as a method for all of science**, from mathematics and physics to economics, the social sciences, media and the humanities, with a strand for high-risk, frontier-moving work and a responsive route for ideas from anyone.
2. **Build AI discovery pipelines around national facilities and data**, including new economic and digital-society data, with universities, including regional ones, as nodes, and routes in for start-ups.
3. **Fund university operations and refresh, not only capital**, and publish what national compute achieves.
4. **Set expectations for verifying AI-generated research**: machine checking where possible, named expert review of what the formal statements claim, and governance standards and model contract clauses for AI providers.
5. **Secure sector-wide access to frontier models on research terms**, including governed access for life-sciences researchers and for university alignment, evaluation and security research, and keep supporting open-weight models.
6. **Make it simple for companies to share university infrastructure**, and fund research engineering careers at scale.

1 Introduction

1.1 Purpose and audience

This report explains why I believe **AI is becoming the core method of scientific research**, what that means for universities, and how policy should evolve to support it. It is written for university leaders, research and innovation directors, developers of innovation campuses, policymakers and research funders, and for founders and investors building science companies. It is intended to be of use in Ireland, the United Kingdom, the European Union and the United States. By "core method" I mean a general-purpose instrument used across the research cycle, as statistics and computing became, not a replacement for experiment, fieldwork or interpretation. By AI I mean the whole family of methods, not only the large language models (LLMs) behind today's agents. It includes domain models such as AlphaFold and AI weather forecasters, which have already transformed their fields. It also includes newer approaches

such as world models, which learn how physical environments behave so that agents and robots can be trained in simulation (Section 5.6).

1.2 My perspective

I am a computer scientist by background and have worked in university research and commercialisation, in the start-up community and as an investor. I founded FeedHenry, acquired by Red Hat in 2014, and was CEO of the Walton Institute (formerly TSSG), a university research institute. I have seen at first hand what it takes to build and run research infrastructure, to turn research into companies and to fund them.

This report sets out my view in my role as Managing Partner of SVV. It draws on recent work with universities on how AI should shape their research strategies and innovation campuses, and on what we see in the companies we back. It sets out a general approach rather than a plan for any one institution.

1.3 Approach and limitations

The report is based on a structured desk review of publicly available sources published up to 5 November 2026, with selected updates to 21 November, including peer-reviewed and preprint literature, company announcements, specialist journalism, and government and funder documents. Throughout, "compute" means accelerated capacity of the kind used to train and run AI models, mostly GPUs today but also other accelerators such as Google's TPUs, not general-purpose servers. "National and international infrastructure" means a country's own systems and those it shares with others, such as the EuroHPC network in Europe.

Four limitations apply. First, many 2026 results are company announcements or preprints not yet peer reviewed. I report them with their caveats. Second, some programme and funding figures depend on future budgets. Third, the forecasts in Section 5 are my judgement, informed by the evidence but not derived from a formal model. Fourth, policy in all four jurisdictions is moving quickly, and some details will have changed by the time this report is read. SVV invests in AI companies that may benefit from the developments described here. That interest, and my use of AI tools in researching and developing this report, are declared in full at the end.

This report is about AI as a method of science rather than AI's wider risks, although the two are hard to separate. The International AI Safety Report 2026 was written by more than 100 experts under Yoshua Bengio and backed by more than 30 countries and international organisations. It groups the risks of general-purpose AI into malicious use, such as cyberattacks and biological risks; malfunctions, including loss of control; and systemic risks such as effects on labour markets. It also records real uncertainty about how fast these will develop [26]. Some of these risks have already materialised: in 2025, Anthropic reported disrupting a cyber-espionage campaign that it attributed, with high confidence, to a Chinese state-sponsored group. Its model had carried out an estimated 80 to 90% of the tactical work against about 30 targets, a handful of them successfully [27]. Experts disagree about how likely the most severe outcomes are, and this report does not try to settle that. The capabilities that make AI useful in research also make it open to misuse, which is why I treat verification, governance and security research as part of the infrastructure (Sections 3.7, 6.8 and 6.12).

1.4 Structure

Section 2 reviews the results frontier AI laboratories produced in 2026, and Section 3 looks at each field in turn, from the life sciences to the humanities. Section 4 describes the wider landscape of companies, open models, national programmes and university compute, and Section 5 sets out my outlook for 2031 and 2036. Section 6 is a playbook for universities, Section 7 covers cost, phasing and funding, Section 8 sets out how policy should evolve, and Section 9 concludes. A separate short briefing for university leaders summarises the recommendations.

2 The state of AI-driven research in 2026

During 2026, frontier AI systems moved from assisting researchers to producing research results, first in mathematics and increasingly in biology and physics, where laboratory validation sets the pace.

2.1 Results from the frontier laboratories

Table 1 summarises the main announcements from OpenAI, Anthropic and Google DeepMind, most recent first. Most are company announcements, so the caveats column should be read alongside the claims. Independent and academic results, including the AutoDNA laboratory published in Nature in November, are covered field by field in Section 3.

Table 1. Selected AI-for-science announcements by frontier laboratories, 2026.

Date (2026)	Laboratory	Announcement	Caveats
20 Nov	Anthropic	Levent Alpöge, an Anthropic mathematician, posted a proof of the Riemann hypothesis, found with Claude, as a Lean formalisation [14], [15], [16], [17]	Accepted by independently written proof checkers; on 21 November the Lean and Mathlib communities certified both the formal statement and the proof [18]; Alpöge's own full account of the mathematics and peer review still to come; Clay lists the hypothesis as unsolved [28]
20 Nov	Anthropic; NYU	Levent Alpöge and Tristan Buckmaster posted a claimed proof that the Navier–Stokes equations can blow up in finite time with no external force, heavily AI-assisted, with a partial Lean formalisation [29]	Unverified: the numerical certificate it relies on had not been released, and that computation sits outside the formalisation; Terence Tao says the strategy is clear but the proof needs more work; Clay lists the problem as active and has not awarded the prize [29], [28]
26 Oct	OpenAI; Anthropic	Each reported, independently, a subexponential-time algorithm for 3-SAT that would refute the Exponential Time Hypothesis (Impagliazzo, Paturi and Zane, 1999–2001), and released Lean proofs [1], [4], [30]	Machine-checked, although OpenAI's Lean file establishes a weaker bound than its paper's main theorem, still enough to refute the hypothesis [31]; Anthropic says its Lean code was written and reviewed only by AI models, that no person is known to have checked it, and that no machine checks that its statement file matches the paper [4], [32]; outside reruns and statement checks have since been reported, but no full review is complete [6]; an independent review of OpenAI's proof, led by Impagliazzo, Paturi and Ryan Williams, says its involvement should not be taken as certification [5]; purely asymptotic, with no practical speed-up, and does not resolve P versus NP [33]
26 Oct	OpenAI	Claimed proofs of Sogge's local smoothing conjecture in all dimensions, which would settle the Kakeya and restriction conjectures, and of the existence of four-dimensional quantum field theories, including pure Yang–Mills, part of a Millennium Prize problem [1], [30]	External review pending; OpenAI claims existence but not a mass gap, so the Millennium Prize problem is not claimed to be solved [34]; the AI-produced physics material is withheld, and a group led by David Gross will write up the work itself; OpenAI says it will not claim the prize [30]
16 Oct	Anthropic, with Harvard	Claude Mythos 5, working in a closed loop with an automated checker, designed circuits for eight standard quantum computing tasks; in seven it reduced both Toffoli-gate and qubit counts below the best published designs, in some cases by an order of magnitude or more [35], [36]	Preprint; designs verified in classical simulation, not run on quantum hardware [35], [36]

Date (2026)	Laboratory	Announcement	Caveats
6 Oct	OpenAI	Released 722 AI-written manuscripts from an internal frontier model [37], reporting progress on 372 problems [38], including a claimed proof of a quasi-Riemann hypothesis, narrowing the strip in which the zeros of the zeta function can lie [2], and manuscripts claiming results on the Unique Games Conjecture and a claimed counterexample to Hadwiger's conjecture itself, a central open problem in graph theory [39]	None peer reviewed, and the Hadwiger claim remains unverified [39]; three manuscripts withdrawn the next day after a sign error and 14 revised; about 42% of top-line results carried Lean formalisations by mid-October [40]; OpenAI said roughly half the results were released unconfirmed [38]; of two quasi-Riemann proofs, the 199-page one (zeros between 1/8 and 7/8) is verified in Lean, while the shorter proof of a weaker bound, which mathematicians are studying, is not; Jared Duker Lichtman called it "the most significant AI-for-math achievement so far, by quite a margin", and Peter Sarnak will not treat it as a theorem until humans have checked the details [2]
23 Sep	Anthropic	About 950 Claude agents searched a DNA database for 21 hours and flagged an enzyme system in bacteriophage DNA not previously described in the literature, with CRISPR-like repeats beside an unusual reverse transcriptase [41], [42], [43]	Preprint [41]; function unknown; the reverse transcriptase was already known, and only its tandem repeat array was new; human scientists did all laboratory work, and Anthropic does not claim a gene-editing tool [42]; outside experts called the result early and incremental [44]; a Copenhagen researcher who had studied the system, unpublished, for about four years while using Claude says his work describes the same system and questions whether it informed the result, and Anthropic says it does not train on user transcripts [45]
8 Sep	OpenAI	About 10,000 coordinated agents produced, in about 88 hours, a proof of finite-time blow-up for the three-dimensional Navier–Stokes equations with a smooth external force, which OpenAI says meets the Clay problem's conditions, formalised in Lean [11], [46]	Not peer reviewed; the Clay Mathematics Institute lists the problem as active and has not awarded the prize [12], [28] and OpenAI says it will not claim the prize [11]; priority and attribution are disputed by NYU's Tristan Buckmaster and Anthropic's Levent Alpöge, who announced related AI-assisted results about 12 hours earlier; both build on a method by the mathematicians Diego Córdoba and Luis Martínez-Zoroa [46]
8 Sep	Google DeepMind	Released the AlphaGenome Atlas: precomputed predictions of the molecular effects of about 9 billion possible single-letter changes in the human genome, free for non-commercial use [47], [48]	Not validated for clinical use; one variant it helped prioritise was confirmed experimentally [47]
6 Sep	OpenAI	Reported reaching its goal of an "automated research intern"; its research organisation now uses about 3.1 agent workdays for every human workday [49]	Self-reported internal metric
4 Sep	Anthropic	Dozens of Claude agents, working largely autonomously for about 11 days, produced the first complete machine-checked proof of Fermat's Last Theorem in Lean, adapting parts of Kevin Buzzard's Imperial College project [50]	Formalises existing mathematics rather than proving new results; about 13 million lines, likely far longer than needed; unreleased internal model
27 Aug	Anthropic	10,000 free or discounted Claude seats for scientists, and AI for Science credits of up to $50,000 per project, which Anthropic plans to extend beyond biology [51]	Biology and chemistry researchers are limited to Opus-class models; access to the most capable models for life-sciences professionals runs through a separate programme set up with the US government [51]

Date (2026)	Laboratory	Announcement	Caveats
1 Aug	OpenAI	An unreleased model reported ten advances in mathematics and theoretical computer science, including solutions to three further Erdős problems [52]	Each with a machine-checked Lean certificate, but the formal statements were recorded as agent-reviewed rather than checked by a human domain expert, and one remains disputed [53]
29 Jul	OpenAI	Free access to frontier models for 10,000 academic researchers, rising to 100,000 through 2027; about 1.3 million people a week already use ChatGPT for advanced science and mathematics [54]	An access programme, not a discovery
20 May	OpenAI	An internal general-purpose reasoning model disproved the Erdős unit distance conjecture (problem posed 1946) [7], [8], [55]	Checked by external mathematicians, who published a human-verified version [8]; Will Sawin soon gave an explicit bound [56]
May	Google DeepMind	Invited any researcher to apply for access to its science tools, including its AI co-scientist, under a new "Gemini for Science" brand [57]	Access by application; the co-scientist is still not publicly available, and commentators noted a gap between the rhetoric and demonstrated results [57]
7 May	Google DeepMind	AlphaEvolve's first-year results included contributions to TPU design, power grid optimisation and a roughly fourfold speed-up of Schrödinger's machine-learned force-field models [58]	Applies where results can be scored automatically
16 Apr	OpenAI	Launched GPT-Rosalind, its first life-sciences model, through a trusted access programme for vetted organisations [59]	Access restricted on biosecurity grounds
5 Feb	OpenAI, with Ginkgo Bioworks	GPT-5 directed six rounds of experiments across more than 36,000 reaction compositions in an automated cloud laboratory, cutting the cost of cell-free protein production by 40% [60]	One protein and one system; experienced human operators were still needed [60]
29 Jan	Google DeepMind	Gemini screened 700 "open" Erdős problems and produced five seemingly novel solutions among the 13 it addressed [61]	"Open" often reflected obscurity rather than difficulty, and some solutions matched existing literature [61]
28 Jan	Google DeepMind	AlphaGenome, which predicts how DNA changes affect gene regulation across sequences up to a million letters long, was published in Nature [62]	Outside experts quoted by AFP called it strong but "far from perfect", limited by its training data [63]

2.2 Four patterns for universities

1. **Progress has been fastest where results can be checked automatically.** Mathematics and algorithm design led because a proof assistant or an automated scorer can check candidate answers cheaply. Where checking requires an experiment, progress moves at the pace of the laboratory. Even in mathematics, cheap machine checking has outpaced the capacity of experts to judge whether formal statements capture the intended problem [53].
2. **Frontier laboratories are building physical laboratories.** Anthropic now operates a wet biology laboratory [43], and OpenAI has run closed-loop experiments in an automated cloud laboratory [60]. Compute alone does not complete the research loop.
3. **Model access is opening up, but only partly.** Free seats and credits for academics are capped, the most capable models are often held back, and restrictions are tightest in the life sciences [51], [54], [59]. To work at the frontier on its own questions and its own data, a university needs compute it controls, alongside distinctive data and automated experimentation.

4. **Claims can run ahead of evidence.** In October 2025, an OpenAI executive announced solutions to ten Erdős problems that turned out to be existing literature [64]. In September 2026, the Anthropic enzyme and OpenAI Navier–Stokes results both drew priority disputes from researchers who were themselves using AI (Table 1) [45], [46]. OpenAI's head of science has said the company is working to give its models "epistemological humility" [65]. Universities have a natural role as independent verifiers.

3 Field by field

AI is arriving unevenly across science, moving fastest where three conditions hold: large, well-curated datasets; a fast and trusted way to check answers; and the ability to act on a prediction, whether by running an experiment or deploying a forecast. Table 2 summarises where each field stands. The fields are grouped by where AI is having most effect, not by a formal classification such as the OECD's Frascati fields: engineering sits with chemistry and materials, agriculture with Earth and climate, and the arts with the humanities. The life sciences receive the most attention below, because they combine the largest data assets, the most industrial activity and the most serious risks.

Table 2. AI as a research method: where each field stands in late 2026.

Field	What AI is doing now	Evidence	Main constraint
Life sciences and medicine	Predicting and designing proteins and antibodies; genome models; AI-designed drugs in late-stage trials; agentic "AI scientists"; automated laboratories	Strong for structure and design; early for clinical outcomes and cell models	Wet-laboratory validation, data on how cells respond, biosecurity
Mathematics	Proving theorems in formal languages; finding new constructions and algorithms; solving open problems	Strong where results are machine-checked; individual landmark claims under review	Expert judgement of what formal results mean
Chemistry, materials and engineering	Predicting stable materials; generative design; robotic synthesis; AI surrogates for engineering simulation	Mixed for materials, where some high-profile AI claims have been corrected after independent checks; proven for surrogates in design work	Synthesis, characterisation and proving novelty; scarce engineering data
Physics, energy and quantum	Plasma control for fusion; particle identification; quantum error correction and circuit design	Strong in specific, well-defined tasks	Access to facilities and hardware
Earth, climate and agriculture	Operational AI weather forecasts; Earth-system foundation models; climate emulators; precision agriculture and dairy	Strong for weather; emerging for climate and agriculture	Extrapolating beyond observed conditions
Social sciences and humanities	Simulated survey respondents and economic agents; reading damaged ancient texts; evidence synthesis	Strong for text recovery and screening; weak for replacing human subjects	Validity against real human data
Computer science and AI	Coding agents writing much new code; automated research on computational problems; alignment, interpretability and security research	Strong for coding; early for automated research	Compute and model access; verifying code written by agents

3.1 Life sciences and medicine

The life sciences are where AI as a scientific method was first recognised at the highest level. The 2024 Nobel Prize in Chemistry went half to David Baker for computational protein design and half to Demis Hassabis and John Jumper for protein structure prediction [66]. By then more than two million people in 190 countries had used AlphaFold 2 [67].

PROTEIN STRUCTURE AND DESIGN

Structure prediction has extended from single proteins to complexes. AlphaFold 3 predicts how proteins interact with DNA, RNA and small molecules. Its code was released for academic use in November 2024 [68]. Open alternatives followed quickly. Boltz-2, released by MIT and Recursion under an open licence, predicts binding strength as well as structure. On a standard benchmark it approaches physics-based simulation (a correlation of 0.62 against 0.72) while running roughly 1,000 times faster [69]. OpenFold3, developed under the OpenFold Consortium, whose members include Novo Nordisk and Bayer Crop Science, is released under an open licence that permits commercial use [70].

Design has advanced as fast as prediction. In 2025, Baker's group reported the first antibodies designed entirely by computer against targets with no known binders, some approaching the strength of clinical antibodies [19]; designed proteins that protected mice from lethal snake-venom toxins [71]; and designed enzymes with complex active sites whose structures matched the computational models to within an ångström [72]. AI-engineered antibodies are moving into late-stage trials: Generate:Biomedicines has announced global Phase 3 trials of an AI-engineered antibody for severe asthma in about 1,600 patients [73]. Not every programme advances: Absci took its AI-designed antibody for inflammatory bowel disease into a Phase 1 trial [74], but later chose not to start later-stage trials itself [75]. Design from plain language is arriving too: in October 2026, Nature published Pinal, a model from a team led by Westlake University that designs proteins from descriptions of their function. Its designs showed the intended activity in all four classes tested, although most screening was computational [9], [76].

One study offers an important caution: when four leading co-folding models, which predict a drug and its protein target together, were tested with physically implausible changes to binding sites, they often still placed the drug molecule in the original pocket. This suggests that they rely on learned patterns rather than physics. The authors advise checking predictions with physics-based methods or experiments before relying on them in drug discovery [77].

GENOME AND CELL MODELS

Language models are now being trained on DNA itself. Evo 2, developed by the Arc Institute and NVIDIA with researchers from Stanford, UC Berkeley and UCSF, was trained on about 9 trillion DNA base pairs from all domains of life and released with its weights, code and training data [78], [79]. In its Nature paper, it predicted the effects of variants in the BRCA1 breast cancer gene without task-specific training, and a simple classifier built on it separated harmful variants from others with high accuracy [79]. Google DeepMind's AlphaGenome, which predicts how DNA changes affect gene regulation, and its atlas of predictions for about 9 billion variants carry the caveats noted in Table 1 [47], [48], [62], [63]. In August 2026, Stanford and Arc researchers reported in Science that genome models had designed entire bacteriophage genomes: of nearly 300 synthesised, 16 produced working viruses, and a cocktail of them overcame bacterial resistance to the virus they were modelled on [80].

The next goal is a "virtual cell" that predicts how cells respond to drugs or genetic changes. The Chan Zuckerberg Initiative has consolidated its science under Biohub, acquired EvolutionaryScale, the developer of the ESM protein models [81], and plans to scale its computing to 10,000 GPUs for biological

research [82]. In October 2026, a cross-sector collaboration announced nearly $1.8 billion for AI-ready biological data, which the collaboration headlines as nearly $2 billion [83]. About $1.3 billion (€1.2 billion) is new money: more than $500 million from the US Department of Energy through the Genesis Mission, $500 million from Biohub, and $300 million from Google DeepMind, Isomorphic Labs and Meta. The rest is coordination by the National Institutes of Health (NIH) of datasets from more than $500 million of earlier federal investment [83].

The focus on data follows from recent results. In the first Virtual Cell Challenge, the organisers reported that entries were "not yet consistently outperforming naive baselines across all metrics" [22]. An independent study in Nature Methods found that deep-learning models did not beat simple linear baselines at predicting the effects of genetic perturbations [84]. In cell modelling, the limiting factor is now data more than method.

DRUG DISCOVERY AND THE CLINIC

Among small-molecule drugs whose target and molecule were both identified with AI, Insilico Medicine's rentosertib, for idiopathic pulmonary fibrosis, is one of the most advanced. A 71-patient Phase 2a trial in China, published in Nature Medicine, was designed primarily to test safety. At the highest dose, forced vital capacity, a measure of lung function, rose by a mean of 98.4 mL over 12 weeks, while it fell by 20.3 mL on placebo [85]. A Phase 3 trial in about 320 patients in China began in July 2026 [86]. Isomorphic Labs, Google DeepMind's drug discovery company, has moved its expected first clinical trial from the end of 2025 to the end of 2026 [87].

The evidence on success rates is still thin. The most widely cited analysis found that AI-discovered molecules had an 80 to 90% success rate in Phase 1 trials, well above historical averages, but about 40% in Phase 2, in line with them, on a limited sample [88]. The first generation of AI-designed drugs largely failed in the clinic [89], and Recursion, the largest listed AI-native drug company, cut about 20% of its staff in 2025 and narrowed its pipeline [90]. AI is shortening early discovery. Whether it improves the odds of a drug working in patients is not yet proven.

AI SCIENTISTS AND AUTOMATED LABORATORIES

Several groups have built agentic systems that generate hypotheses, analyse data and propose experiments. Google's AI co-scientist proposed drug-repurposing candidates for leukaemia that were then confirmed in the laboratory, and independently arrived at an unpublished mechanism of bacterial gene transfer [91]. Stanford's Virtual Lab, a team of AI agents with a "critic" agent to check their work, designed 92 nanobodies against SARS-CoV-2, two of which also bound recent variants of the virus [92]. Kosmos, from FutureHouse, reads about 1,500 papers and runs about 42,000 lines of analysis code in a single run. Independent scientists judged 79.4% of the statements in its reports to be accurate [93], which is impressive and also means that about one in five was not. Commenting on Anthropic's enzyme result (Table 1), one biologist noted that its agents were guided by senior domain experts, and that such systems are "really good at grunt work when orchestrated by experts in the field" [44].

Automation is also closing the loop between prediction and experiment: a self-driving laboratory at the University of Illinois engineered enzymes with up to 26-fold higher activity with little human intervention [94], and OpenAI has run closed-loop experiments with Ginkgo (Table 1) [60]. In November, Nature published AutoDNA, an automated laboratory in which people were needed only to maintain the hardware (Section 6.7) [3]. Companies such as Lila Sciences are building "AI Science Factories" on this model at industrial scale [95].

CLINICAL RESEARCH

AI is also changing how clinical evidence is produced. In the randomised MASAI trial in Sweden, AI-supported mammography screening detected 29% more cancers and cut radiologists' reading workload by 44% without increasing false positives [96]. Delphi-2M, a generative model trained on about 400,000 UK Biobank participants, estimates risk for more than 1,000 diseases and was validated on 1.9 million people in Denmark [97]. Regulators are responding: the US Food and Drug Administration has issued draft guidance on AI in drug development, though it explicitly excludes drug discovery [98].

BIOSECURITY AND ACCESS

The same capabilities raise real risks, and access to the most capable tools now depends on governance. Anthropic introduced stronger safeguards for biological and chemical misuse in 2025, with targeted exemptions for vetted researchers with legitimate dual-use applications [99]. OpenAI offers its life-sciences model only through a trusted access programme [59]. The US AI Action Plan recommends requiring federally funded institutions to use nucleic acid synthesis providers that screen orders and verify customers [100]. The UK has published screening guidance for synthetic nucleic acids [101], and the EU's proposed Biotech Act would require benchtop DNA synthesisers to include built-in screening [102]. For a university, biosecurity governance is becoming a condition of access to frontier tools in biology.

GOVERNMENT PRIORITIES

Governments see the life sciences as the first test of AI for science. The UK's AI for Science Strategy makes drug discovery its first mission, aiming for trial-ready drugs within 100 days by 2030 [103], and its OpenBind consortium released 699 X-ray structures of compounds bound to a viral protein in May 2026 [104]. The US Genesis Mission names biotechnology among its priority domains (Table 4) [105].

For universities, the lesson is that the life sciences reward the assets that are hardest to copy: patient cohorts, biobanks, hospital partnerships, wet laboratories and experimental data. A university with these assets, and the compute and governance to use them, is well placed, while one without them will find it hard to compete on models alone.

3.2 Mathematics

Mathematics moved from benchmark to research frontier in under two years. At the 2025 International Mathematical Olympiad, an advanced version of Google's Gemini Deep Think scored 35 of 42, a gold-medal standard certified by the IMO, working in natural language within the time limit [106]. DeepMind's AlphaProof, which writes proofs in the formal language Lean, was published in Nature the same year [107]. AlphaEvolve, which pairs a language model with automated evaluation, found a way to multiply 4×4 complex matrices with 48 scalar multiplications, one fewer than Strassen's 1969 method allows [108]. In a study co-authored by Terence Tao, it matched or improved the best-known constructions across 67 problems [109].

The Erdős problems show how claims have moved from overstatement to substantive results. After the overclaim of October 2025 [64], the Rényi Institute credited ChatGPT with solving five Erdős problems in January 2026 [110]. A systematic DeepMind study warned that "open" often reflected obscurity rather than difficulty (Table 1) [61]. The May 2026 unit-distance result was the first that many leading mathematicians treated as a major advance [111], and human mathematicians soon published a verified version and an explicit bound [8], [56].

The infrastructure for machine-checked mathematics is growing. The Equational Theories Project, led by Tao, settled more than 22 million implications between algebraic laws by combining human and auto-

mated proofs, all checked in Lean [112]. Philanthropy has followed: the AI for Math Fund, supported by XTX Markets, has reached $31.5 million [113], and DARPA's expMath programme, which held its kick-off in March 2026, aims to build an AI co-author for mathematics [114].

Formal verification does not settle everything: the Navier–Stokes proof was checked in Lean, yet whether its formal statement captures the problem, and who deserves credit, remain in question (Table 1) [46]. A preprint argues that Lean verification of an AI translation does not guarantee a correct natural-language proof, and reports mismatches in examples including that proof [13]. The October 2026 claims to refute the Exponential Time Hypothesis are the clearest example. Both proofs passed machine checks. When they were announced, no person was known to have checked Anthropic's Lean code [4]; outside volunteers have since rerun both proofs, and one has read both formal statements by hand, but no full independent review has been completed [6]. Two of the hypothesis's originators, who are leading an independent review of OpenAI's proof, say their involvement should not be taken as certification [5]. There is also a question of openness: the Advisory Group on Mathematics and Artificial Intelligence, an independent group of mathematicians hosted by the Institute for Advanced Study, has asked AI laboratories to stop testing advanced mathematical problems on proprietary models [115], [116].

The proof of the Riemann hypothesis in November shows the same division of labour. Machine checking was fast and independent, but confirming that the formal statement says what mathematicians mean was human work, done by the Lean and Mathlib experts who certified it on 21 November [18].

3.3 Chemistry, materials and engineering

Materials science shows both the promise and the risk of overclaiming. Google DeepMind's GNoME predicted about 2.2 million crystal structures, around 380,000 of them stable, and contributed nearly 400,000 compounds to Berkeley Lab's Materials Project [117]. Independent chemists found "scant evidence" of compounds that were at once novel, credible and useful, although DeepMind stands by its claims [118]. Berkeley's autonomous A-Lab originally reported synthesising 41 new compounds in 17 days. After outside researchers questioned whether the products were new, the Nature paper was corrected, and it now reports 36 of 57 targets synthesised [119], [120]. Three years on, A-Lab runs at roughly 100 times human iteration speed [121].

Generative design is maturing. Microsoft's MatterGen designed a material for a target bulk modulus (resistance to compression) of 200 GPa. When synthesised, it measured 169 GPa, within 20% of the target [122]. Universities have been pioneers in automation: at the University of Liverpool, mobile robots share standard laboratory equipment and decide which reactions to run next [123]. The University of Toronto received C$200 million for self-driving laboratories, aiming to cut materials development from about 20 years to one [124]. Capital is following, as Periodic Labs' $300 million seed round for autonomous materials laboratories shows (Table 3) [125].

Engineering shows a similar pattern in design. AI surrogate models now stand in for expensive simulations: at Leonardo Helicopters, a geometric deep-learning model predicts ground-level rotor noise about 540 times faster than conventional computational fluid dynamics. In casting, an AI model predicts porosity in aluminium parts almost instantly where a simulation took a day [126]. Engineers describe AI as an accelerator for physics rather than a replacement, held back by scarce, proprietary data, sometimes as few as 20 simulations [126].

AI can generate candidates and designs far faster than they can be confirmed, so the bottleneck lies in synthesis, characterisation, testing and judgement of novelty, all of which universities do well.

3.4 Physics, energy and quantum

In physics, AI is strongest on specific, well-defined tasks. Google DeepMind is working with Commonwealth Fusion Systems to optimise plasma control for the SPARC fusion device using reinforcement learning and an open-source simulator [127]. At CERN, the ATLAS experiment's transformer-based tagger improved the rejection of charm-quark jets 3.5-fold and of light-quark jets 1.8-fold over its predecessor, at 70% b-jet efficiency [128].

Quantum technology is where the link between AI and a hard engineering problem is now clearest. AI is accelerating quantum research in at least four ways:

- **Error correction.** Google DeepMind's AlphaQubit 2 decodes errors in real time, in under a microsecond per cycle, on surface codes up to distance 11 and colour codes up to distance 9 [129]. An AI system has also proposed error-correcting codes that outperform leading designs, so far in a preprint [130].
- **Circuit design.** Claude designed quantum circuits that beat the best published designs in seven of eight tasks (Table 1) [35].
- **Hardware design.** IBM's AI agents optimised the hardware design of an error-correction decoder, cutting its estimated latency 52-fold for about $26 of computing with open models. These are design-tool estimates on a reduced problem, measured against an unoptimised starting design rather than IBM's hand-tuned decoder, not hardware measurements [131], [132].
- **Laboratory work.** Microsoft credited an AI agent with automating measurements for its Majorana 2 chip, whose qubits it says hold their state 1,000 times longer than its first generation. Some physicists say Microsoft has not released enough data to verify its claims [133], [134].

Quantum computing itself is moving towards engineering: Quantinuum targets a fault-tolerant computer by 2029 [135], and IBM and the University of Chicago have claimed a verifiable quantum advantage in preprints awaiting assessment [136], [137]. Google's Quantum Echoes result points to molecular simulation as a leading candidate for early useful advantage [138]. It is unlikely to replace classical computing within the decade, but likely to join the research stack as a specialised accelerator for simulation. A university quantum group is best connected to the AI stack, with simulators on the institution's GPU cluster and access to physical machines through national and commercial services, rather than run as a separate facility.

3.5 Earth, climate and agriculture

Weather forecasting is the clearest example of AI moving from research into operational use. Google DeepMind's GraphCast outperformed the leading physics-based forecast on more than 90% of test measures [139], and its probabilistic successor GenCast was more accurate than the European Centre for Medium-Range Weather Forecasts (ECMWF) ensemble on 97.2% of targets [140]. In 2025 ECMWF made its own AI forecasts operational, alongside its physics-based system, using about 1,000 times less energy per forecast [20], [141]. Microsoft's Aurora, trained on more than a million hours of geophysical data, beats operational forecasts for air quality, ocean waves and cyclone tracks [142].

Climate projection is harder. Hybrid models such as NeuralGCM run stably for decades at a fraction of the computing cost, but do not reliably extrapolate to climates very different from those they were trained on [143]. Emulators such as ACE2 can simulate about 1,500 years a day but respond imperfectly to some forcings [144]. ECMWF's AI ensemble still depends on physics-based data assimilation [141]. For now, AI speeds up physics-based climate modelling rather than replacing it.

Agriculture and food draw on the same kinds of data. Ireland's VistaMilk centre, hosted by Teagasc with 13 research institutions and more than 50 industry partners, applies AI, sensors and mathematical modelling from soil and pasture to animal health and dairy processing, and began a €26.2 million second phase in 2025 [145]. For many regional universities, agri-food is a natural domain for the data-led models described in Section 6.5.

3.6 Social sciences and humanities

In the social sciences, AI offers a new instrument, though it is tempting to use it in place of human respondents. Researchers showed in 2023 that language models conditioned on real respondents' backgrounds could reproduce the response patterns of survey subgroups, which they called "silicon samples" [146]. Later work found that such samples showed too little variation, that about half of the relationships estimated from them differed significantly from those in human data and about a third of those had the wrong sign, and that results shifted with prompt wording and model updates [147]. A 2026 preprint from Prolific found that giving simulated respondents demographic personas roughly tripled the error in the distribution of answers, compared with asking the model directly for an aggregate distribution. In the same preprint, the method that reproduced the distribution worst was nonetheless among the best at predicting the most common answer [148]. Centaur, a model trained on 160 psychology experiments with more than 60,000 participants, predicts human choices better than existing cognitive models [149]. Critics, however, argue that it does not work like a human mind [150] and that it may be memorising its training data [151]. In economics, language models are now widely used as research assistants [152] and as simulated economic agents in classic experiments [153].

In the humanities, AI is opening sources that were previously unreadable. The Vesuvius Challenge, combining synchrotron imaging with AI, has recovered more than 70 columns of text from one carbonised Herculaneum scroll and around 20 from another, possibly an unknown work of Stoic ethics [21]. DeepMind's Aeneas, trained on about 150,000 Latin inscriptions, helped historians generate research ideas for 90% of the inscriptions they studied [154].

Evidence synthesis is changing fastest. Cochrane, the Campbell Collaboration, JBI and the Collaboration for Environmental Evidence have jointly accepted AI in systematic reviews with human oversight and full reporting [155]. In a preprint, an AI workflow reproduced and updated 12 Cochrane reviews in two days, with higher screening sensitivity than human reviewers, and its updates changed the conclusions of three [156].

The lesson for these fields is that AI works best as an instrument checked against human data and expert judgement, not as a substitute for human participants. Where AI is used with human-subject data, or simulated respondents stand in for people, research ethics committees should review it (Section 6.12).

3.7 Computer science and AI itself

AI is itself a product of computer science, and the field is changing faster than any other. Much of the groundwork came from universities: the 2024 Nobel Prize in Physics recognised the foundations of machine learning [157]. The frontier has since moved to industry. In 2025, industry produced 93 of the notable AI models tracked by Stanford's AI Index and academia 2, although academia still produced most AI publications (68% in 2024) [158]. The gap is one of compute rather than ideas, and it shapes what university AI research can now attempt.

How software is built is changing first. Coding agents now write much of the code at leading companies: Google says three-quarters of its new code is generated by AI and then reviewed by engineers [159]. As

agents produce code faster than people can read it, the bottleneck moves from writing code to specifying what it should do and proving that it works. Across GitHub, more than 40% of pull requests started with an agent's first draft by September 2026, up from 17% in March. Of agent-involved pull requests, 90% were merged within two weeks, against about 86% of others [23]. The gains in code written are large; the gains in software shipped are smaller (Section 5.1). Computer science is the first field in which the verification problem described in this report has appeared in practice. It affects universities directly: much research is software, research engineers will increasingly specify and verify code rather than write it, and computer science teaching will have to change with them. An end-to-end automated research system has already produced a paper that passed the acceptance threshold at a 2025 workshop, although it remains limited to computational work and has hallucinated citations [160], [161].

The research that matters most for AI's own future is open to universities too. Alignment, interpretability, evaluation and security are active fields in which the frontier laboratories do much of the work, but universities bring independence, open publication, long time horizons and evaluation free of commercial interest. That matters more now that AI may be speeding up its own development. METR judges that most of Anthropic's recent acceleration is likely due to AI [10], and Anthropic's amended IPO filing says it may imminently cross its own threshold for AI-driven acceleration [162]. Universities can still do frontier work here: in November, MIT researchers reported that OpenAI's GPT-6 Astra can read and directly edit smaller models' weights, raising one small open model's mathematics score from 8% to 48%, although it could not detect a sycophantic model from its weights alone [163]. Funders and laboratories are reaching out to universities. The Alignment Project is a coalition led by the UK AI Security Institute with partners including UKRI, ARIA, the Canadian AI Safety Institute, OpenAI, Anthropic and Microsoft. In its first round it awarded more than £27 million to more than 60 projects, with compute and API credits alongside grants [164]. Anthropic, OpenAI and Google run residencies and visiting-researcher programmes [165], [166]. The constraint for universities is access: research on frontier systems needs model access and compute that universities rarely hold, another reason for the agreements and shared compute in Section 6.4. My forthcoming companion report, AI at the Core: Who Checks the AI?, sets out these research areas and the routes in for universities.

3.8 Across fields

Researchers are adopting AI faster than their institutions are supporting it. In Wiley's survey, AI use among researchers rose from 57% in 2024 to 84% in 2025, while the share of use cases in which researchers believed AI outperformed humans fell from over half to under a third [167]. Elsevier found that only 22% of researchers consider AI tools trustworthy, 27% feel adequately trained and 32% think institutional governance is good [168]. In Springer Nature's 2026 survey of more than 10,000 researchers, only 11% had institutional funding for their main AI tool [169]. Finance is part of the reason: the Office for Students expects 42.7% of English higher education providers to report a deficit in 2025–26 [170].

Take-up is also uneven across fields. In Anthropic's analysis of about 207,000 research sessions on Claude, use was highest in the formal sciences and lower in laboratory fields such as chemistry and medicine, though these are one company's consumer data [171]. OpenAI's study of how 1,819 scientists used ChatGPT, drawn from more than 580,000 messages, found that most use went on reading, writing, analysis and code, with very few requests to run research end to end. Surveyed scientists estimated that the same work would have taken 1.9 times as long without AI. Their most common difficulty was checking and correcting AI output. The data come from one company and a self-selected sample [172].

Some of the earliest institutional gains may come from the work around research rather than research itself. Much of its cost lies in administration, governance and maintaining the facilities in which it hap-

pens: of the more than $35 billion NIH spent on grants in 2023, about $9 billion went on indirect costs [173], and US principal investigators report spending about 44% of their research time on administration [174]. Grant preparation, compliance, ethics review, reporting and facility management are all areas where AI can make universities more efficient now, and the savings can help fund the infrastructure in Section 7.4.

The effects on careers and on science itself pull in different directions. A study of 41.3 million papers found that scientists who use AI publish three times as many papers, receive nearly five times as many citations and become project leaders earlier (an association, not necessarily a causal effect). It also found, however, that AI adoption narrows the range of topics science explores and reduces engagement between scientists [175].

Rigour is a cross-field concern. Data leakage errors have been documented in 294 papers across 17 fields that use machine learning [176], prompting consensus reporting standards [177]. A widely cited 2024 study of AI in a materials laboratory was withdrawn in 2025 at the request of MIT, which said it had no confidence in the data [178]. These failures are reasons to build verification into AI-driven research from the start, rather than to stand aside from it.

4 The landscape

Beyond the frontier laboratories, investment and policy are converging on the same capabilities: compute, open domain models, proprietary experimental data and automated laboratories. Governments are building them nationally and start-ups privately, and a growing number of universities are building their own.

4.1 AI-native science companies

Table 3. Selected AI-native science companies.

Company	Approach	Scale
Lila Sciences (Flagship Pioneering)	Closed-loop "AI Science Factories": AI designs experiments, robotic laboratories run them, and the results retrain the models	About $550M raised by October 2025 [95]; reported in June 2026 to be in talks to raise at about $8.5B before new money [179]
Periodic Labs	Autonomous laboratories for materials science, founded by former OpenAI and Google DeepMind researchers	$300M seed round in 2025 [125]; reported in May 2026 to be in talks at about $7.5B [180]
Isomorphic Labs (Alphabet)	AI drug design engine built on the AlphaFold lineage	Raised $2.1B in May 2026 [181]; first clinical trials expected by the end of 2026 [87]
Edison Scientific (FutureHouse spin-out)	Kosmos, an AI scientist that reads literature and runs analyses	Company-reported estimate from beta users that one run is equivalent to about six months of doctoral or postdoctoral work [182]; see Section 3.1 for an independent accuracy check
Boltz	Public benefit corporation formed by the MIT developers of the open Boltz models	$28M seed round and a Pfizer partnership in January 2026 [183]
Xaira Therapeutics	AI drug discovery built on protein and antibody design models developed in David Baker's laboratory at the University of Washington	Launched with more than $1B in committed capital in April 2024 [184]
CuspAI	Cambridge start-up building AI models that design new materials	$450M Series B in July 2026, with the UK's Sovereign AI fund among the investors [185]

Most of these companies are betting that proprietary experimental data from automated laboratories, rather than internet text, is now the scarce asset. Lila has said that leadership in AI for science will depend on owning the largest automated laboratory, not just the biggest data centre [95].

4.2 Open and domain-specific models

Many of the strongest scientific models are open, and many originate in universities and non-profit institutes: Evo 2 from the Arc Institute and its partners [78], Boltz-2 from MIT and Recursion [69] and OpenFold3 from an academic and industry consortium [70]. AlphaFold 3 is available for academic use only [68]. The US Department of Energy has launched a Genesis Open Models Initiative to build open-weight science models with outside contributors, starting with Genesis-Science-1, being developed with Arcee [186]. Europe has its own open-weight developer in Mistral AI, which presents itself as a sovereign European alternative and raised €3 billion at a valuation of more than €21 billion in September 2026

[187]. Open models matter to universities because they can be adapted to local data and run on infrastructure the institution controls.

4.3 National programmes

Governments in the United States, the United Kingdom and the European Union now treat AI for science as a national priority (Table 4). So do others: in August 2025 China's State Council called for new "AI-driven scientific research paradigms" and large scientific models as part of its "AI+" initiative [188]. This report concentrates on the four jurisdictions most relevant to its readers, and Section 8 compares their approaches in detail.

Table 4. Selected national and European AI-for-science programmes.

Programme	Description
US Genesis Mission	Launched by Executive Order 14363 in November 2025 and led by the Department of Energy (DOE). It is building the American Science and Security Platform, linking national laboratory supercomputers and secure cloud environments, AI agents, domain foundation models, federal and partner datasets, and tools for autonomous experimentation. Priority domains include advanced manufacturing, biotechnology, critical materials, nuclear fission and fusion energy, quantum information science, and semiconductors and microelectronics [105]. The administration announced more than $5 billion in federal commitments in July 2026 [189]. DOE reports a first-year portfolio of 297 projects, including the 278 selected in the first round in July 2026 [190], [191]
US National AI Research Resource (NAIRR)	A National Science Foundation (NSF) programme connecting researchers and educators to compute, data and models, piloted from 2024 with private partners and reaching about 900 research teams and educators. In September 2026 NSF made a $35M, five-year award to the San Diego Supercomputer Center, with the Texas Advanced Computing Center, to run a NAIRR Operations Center [192] as NAIRR transitions to a permanent NSF capability [193]
US NSF Programmable Cloud Laboratories	$380M from NSF, with further support from the Astera Institute, to 20 teams building AI-programmable cloud laboratory nodes, announced in July 2026 as NSF's core contribution to the Genesis challenge on autonomous laboratories [194]; universities, non-profits and companies were eligible to apply [195]
UK AI for Science Strategy and AI Research Resource	A strategy published in November 2025 with up to £137M, a first mission on drug discovery, AI Research Resource (AIRR) allocations of 200,000 to 1 million GPU hours over six months for large projects and a "system takeover" call of up to 1.4 million GPU hours over two weeks, an autonomous-laboratories initiative and a target for UKRI facility data to meet FAIR principles (findable, accessible, interoperable and reusable) by 2030 [103]. The government plans a 20-fold expansion of the AIRR by 2030, alongside up to £750 million for a new national supercomputer at Edinburgh [196]
EU Strategy for AI in Science and RAISE	A strategy adopted in October 2025, creating RAISE, a virtual institute pooling compute, data, talent and funding, with a €107M pilot and €600M for researcher access to AI Gigafactories [197]. EuroHPC, which oversees 19 AI Factories and 13 AI Factory Antennas, now has a mandate that also covers AI Gigafactories and quantum technologies [198]. In July 2026 it opened a call for up to seven AI Gigafactories, to be built by consortia of companies, public bodies and investors, with public funding as an anchor customer, and expected to draw more than €20 billion of private investment [199]. In October 2026, at the AI in Science Summit in Dublin, the Commission launched a pledging scheme through which companies offer compute, data and models to researchers [200]

These programmes share a common shape: national compute, national data, priority domains, and access schemes for researchers and companies. Universities take part prominently, as award holders, hosts of national systems and trainers of talent. What remains unclear is how university-hosted AI infrastructure is sustained once built (Section 8).

4.4 University and community compute

Leading universities increasingly own GPU capacity rather than relying only on national facilities. Harvard's Kempner Institute operates about 1,100 GPUs, and the university is considering a further $25 million, not yet finalised, which the institute's co-director estimates could roughly double central capacity [201]. Princeton funded a 300-GPU cluster largely from its endowment so that academics can set their own AI research agenda [202]. Smaller institutions are acting at their own scale: Mid Sweden University expects its EU-funded GPU cluster to strengthen collaboration with industry [203]. In the UK, two universities host national AI systems: Bristol runs Isambard-AI and Cambridge runs Zenith, both part of the AIRR, which offers researchers allocations and UK small firms a rapid access route [103], [204].

Cambridge-based CommonAI CIC offers a different model, one example of an approach that may suit an innovation campus [205]. It is a non-profit membership organisation through which start-ups, established companies, public bodies and universities share cost-effective compute, a permanent team of engineers and researchers, and a digital commons under a protected IP framework. Members join for the shared infrastructure, and their membership helps to sustain it.

A UK precedent shows how national computing capability can anchor a campus outside a major city. Sci-Tech Daresbury in Cheshire, a joint venture between the developer Langtree, the Science and Technology Facilities Council and Halton Borough Council, hosts the Hartree Centre, a national high-performance computing and AI facility. In 2017 a collaborative building opened there, housing IBM Research staff alongside the Hartree Centre and the University of Liverpool to develop computing tools for companies [206]. A survey of 97 campus companies, published in 2018, found that 78% collaborated with at least one other organisation. Of those, 54% used the Science and Technology Facilities Council's services or expertise, up from 41% in 2016, and overall 67% collaborated with a university or the Council [207]. The campus welcomed its 150th tenant in December 2019 [208] and reported more than 150 businesses on site in 2025 [209]. I have not found a more recent survey of how campus companies use the national facilities.

There are comparable models elsewhere. The National Center for Supercomputing Applications at the University of Illinois runs an industry partner programme. In 2018 it had 35 industrial partners, had worked with more than a third of the Fortune 50, and offered smaller firms a dedicated entry-level system [210]. In Finland, the LUMI supercomputer in Kajaani has anchored a regional data-centre ecosystem [211]. The LUMI AI Factory has offered services for start-ups, SMEs and academics on LUMI since October 2025 [212]. A dedicated €387.8 million LUMI-AI system, contracted in August 2026, is due to open to users in 2027 [213], and XTX Markets plans to invest more than €1 billion in data centres in the town [214].

Evidence of what such infrastructure produces is still thin, and mostly reported by the institutions themselves. The Hartree Centre reports supporting 844 businesses across its programmes, and 179 collaborative research and development projects with 80 partners under its £210 million Hartree National Centre for Digital Innovation [215]. The University of Liverpool's partnership with Unilever, anchored in the £81 million Materials Innovation Factory, is reported to have contributed to 350 Unilever patent filings and seven university spin-outs between 2015 and 2024 [216]. Toronto's Acceleration Consortium attracted matching commitments of about C$300 million from 30 partners alongside its C$200 million federal grant, and has launched four start-ups [217], [218]. These are programme-wide figures, not controlled evaluations of what AI infrastructure adds.

5 Outlook: research in 2031 and 2036

These forecasts run to 2031 and 2036 because universities plan infrastructure over five to ten years. Some 2031 expectations are already partly visible: mathematics moved from benchmarks to research results in under two years (Section 3.2), and agents already draft a large share of new code (Section 3.7). Universities, however, move at the pace of funding, approvals and laboratory work rather than of model releases.

5.1 Research in 2031

By 2031, a typical research group is likely to be a small human team directing a large number of AI agents. In most computational work, agent hours will exceed human hours, and the scarce inputs will be compute, experiments, distinctive data and good questions.

This is the forecast I hold with most confidence. OpenAI already reports more agent workdays than human workdays inside its own research organisation, on its own measurements [49]. Anthropic's own index has Claude leading 26% of its AI research work in August 2026, weighted by staff time, up from under 1% in February [219]. OpenAI estimates that its agents succeeded unaided on 46% of the two-to-four-hour tasks its researchers delegated to them in October, up from 23% in March [220]. These are the laboratories' own measures, and none has yet been checked independently. A separate, independent assessment points the same way: METR, with access to Anthropic's internal data, estimates that Anthropic is now improving its models on well-specified tasks about twice as fast as in 2024–25 (plausibly 1.3 to 3 times), mostly because of AI, although Anthropic redacted much of the evidence [10]. When METR included frontier models from other laboratories, its estimate fell to about 1.4 times [10]. There is not yet comparable evidence from universities. Meanwhile, in Wiley's 2025 survey, 84% of researchers surveyed used AI in some form [167]. The pattern is likely to spread through universities over five years, much as AI coding tools spread through software teams in 2025. Adoption is not the same as productivity, and the evidence has moved quickly. An early randomised trial, of 16 developers using early-2025 tools, found them 19% slower [221]. Its authors now say current tools probably speed developers up, but that their trial design can no longer measure the gain reliably [222]. The largest study so far uses field data on more than 500,000 GitHub developers. With autonomous coding agents, commits rose by a cumulative 240%, but software releases by only about 30%, and new apps in software marketplaces rose sharply while total usage did not [223]. The gains are real, but they shrink as work moves from writing code to shipping it, because people still have to review and integrate it. There is also a cost to learning: in a randomised trial, developers who learned a new library with AI scored 17 percentage points lower on a later test of their understanding [224]. My forecast concerns the share of work done by agents, not a guaranteed gain in output per researcher. Randomised field experiments with 4,867 developers found 26% more completed tasks among developers using an AI coding assistant [225].

5.2 Changes in research practice

- **Agents do the groundwork.** Literature synthesis, data analysis, simulation set-up, coding and first drafts are delegated. A doctoral student supervises systems such as Kosmos [93] rather than spending months on each analysis, and checks their output with care. Doctoral training will need to guard against deskilling and against a narrowing of research questions, which AI adoption has already been associated with [175].

- **Machine-checked results spread, with expert review alongside.** In mathematics and theoretical computer science, formal proofs become common, together with named expert review of what the formal statements claim [53]. Other fields adopt auditable agent records linking each finding to the code, data and sources that produced it, and reporting standards of the kind already agreed for machine-learning science and evidence synthesis [155], [177].
- **Grants budget compute like equipment.** Token and GPU budgets become standard line items, and compute per researcher becomes a recruitment factor, as laboratory space and start-up packages are today.
- **Domain models are adapted locally.** Research groups fine-tune open models such as Evo 2 [78] or Boltz-2 [69] on their own data. Institutions with distinctive clinical, agricultural, marine, environmental or manufacturing data gain disproportionately.

5.3 Shifting bottlenecks

- **Experimentation.** Ideas become cheap, but testing them remains expensive. Self-driving laboratories spread in chemistry, materials and synthetic biology, first in companies and well-funded institutes [121].
- **Data.** In fields such as cell biology, models are ahead of the data needed to train them (Section 3.1). Generating high-quality experimental data becomes a strategic activity in its own right.
- **Verification and judgement.** Results may outpace the capacity to check them. Peer review, reproducibility and authorship norms come under strain, and choosing and checking the right question becomes the central human skill.
- **Governance of dual-use research.** With providers already restricting their most capable models for life-sciences work, universities will need biosecurity and data controls to access the best tools [51], [59].

5.4 Organisational change

Laboratory roles are likely to shift towards research engineers, automation technicians and AI operations staff, with fewer junior hands-on roles. The boundary between universities and companies blurs further. The largest closed-loop laboratories are likely to sit in companies such as Lila Sciences and Periodic Labs, and universities will partner with them for access while competing on talent and data. Institutions without a compute and data strategy risk losing their strongest researchers to those that have one. The same logic applies to companies: research-intensive firms are likely to choose where to locate partly on access to compute, data and AI engineers, which matters for every university with an innovation campus (Section 6.9).

5.5 Research in 2036

By 2036, in fields where experiments can be automated and results checked, the loop from hypothesis to validated result is likely to run with little human intervention. Researchers will set goals, judge significance and carry accountability.

- **Research capacity decouples from headcount.** Output scales with compute, automated laboratory throughput and the capacity to verify results. A mid-sized institution with the right infrastructure could advance discovery faster than a much larger one without it.

- **Discovery rates rise unevenly.** The largest gains are expected in mathematics, protein and drug design, materials, chemistry and simulation-intensive fields (Section 3). Fields that depend on human participants, field sites or slow biological processes, such as clinical trials, ecology and ageing, accelerate less, because the physical world sets the pace.
- **Automated "science factories" become a category of research infrastructure.** Facilities combining compute, robotics and data will be built by companies, governments and some university consortia, as synchrotrons and genome centres were in earlier periods.
- **National facilities become discovery engines.** Synchrotrons, health and genomic data resources, Earth observation and climate archives are connected to compute, domain models and automated experiments. This lets universities, public programmes and companies join at the right layer rather than each starting from scratch. Early steps are visible: Argonne is linking its upgraded Advanced Photon Source to national supercomputers for near-real-time, AI-assisted analysis, as part of DOE's Integrated Research Infrastructure effort [226]. Diamond Light Source leads OpenBind, which generates open protein–ligand binding data at its XChem facility and trains AI models on it using Isambard-AI [227].
- **The value of the university shifts.** As producing papers becomes cheaper, universities will be valued for what is hard to replicate: distinctive instruments, field sites, patient cohorts and regional industry data; trusted, independent verification; ethical and safety governance; and the education of people able to direct machine research.

5.6 Uncertainties

I hold the five-year view with reasonable confidence for computational fields and moderate confidence for experimental ones. The ten-year view indicates direction rather than timing. Eight factors could slow it or change its shape.

1. **Laboratory automation lags.** General-purpose laboratory robotics is considerably harder than software, even though specialised systems such as the AutoDNA laboratory (Section 6.7) show what is already possible. If it stalls, experimental science remains human-paced and the gains concentrate in theory and computation.
2. **Specialist models generalise unevenly.** General-purpose frontier models now produce results that were not in their training data, from new proofs to new algorithms (Sections 2 and 3.2). Models trained on narrow scientific data are a different case. Protein–ligand co-folding models still misplace ligands at unfamiliar binding sites [77], [228]. Models that predict the effects of genetic perturbations still struggle to beat simple linear baselines, and they score lower under stricter tests [84], [229]. If that persists, the binding constraint in these fields will be data, not models.
3. **Adjudication becomes the constraint.** Formal proof checking may become nearly free, but deciding which results are correct, important and correctly stated does not. This has been described as verification abundance alongside adjudication scarcity [53].
4. **Compute and energy constrain supply.** If GPU, memory and power availability remain tight, access concentrates in a small number of companies and countries. AI demand for high-bandwidth memory is already tightening memory supply and pushing prices higher [230].
5. **Access to models narrows.** Dual-use concerns, export controls or commercial strategy could restrict the most capable models to approved users, as is already partly the case in the life sciences [51]. In June 2026, a US directive led Anthropic to withdraw its two most capable models from all customers

for about three weeks (Section 6.4) [231]. Continuity of access is now a question of national resilience as well as research policy.

6. **The technology shifts.** Expectations are reset every few months, and the next wave may not look like today's. World models, from Google DeepMind's Genie 3 to NVIDIA's open Cosmos models [232], [233], are attracting large investment: AMI Labs, founded by Yann LeCun, raised $1.03 billion in March 2026 to build them [234]. Accelerators other than NVIDIA's GPUs are also maturing (Section 7.6). Institutions that commit large sums to one hardware generation or one kind of model risk being left with the wrong infrastructure.
7. **Claims outrun reality.** If prominent AI discoveries prove less novel than announced [45], [64], funders may withdraw before the substantive gains arrive.
8. **Private laboratories outrun universities.** AI may be beginning to speed up its own development (Sections 3.7 and 5.1). Whether that compounds is open. One calibrated economic model finds the feedback loops strengthening but not yet self-sustaining [235]. In a review by the Forecasting Research Institute, an ensemble of AI models put the chance of a "fast year" of AI progress (a rise of 30 points or more on the Epoch Capabilities Index within twelve months) by 2030 at about 40%, and published central estimates for the full automation of AI research ranged from 2027 to 2052 [236]. If that compounds, companies with far more capital and compute could run largely autonomous research at a scale no university can match, while AI tutors erode part of universities' teaching role. Universities could then be left with research that has little direct commercial value but needs costly physical experiments.

I do not think any of these factors would reverse the direction of travel, but each could change its pace and form. Together they strengthen the case for building on owned infrastructure, distinctive data and rigorous verification, bought in stages, rather than depending entirely on externally controlled models.

6 The university playbook

6.1 Build a stack, own what differentiates

A university and its innovation campus are best served by building an AI research stack in layers. That means owning the parts that create lasting advantage (a right-sized compute cluster, data, domain models, a closed-loop laboratory, verification and people) and renting or federating the rest, without training frontier models of their own. Figure 1 summarises the structure.

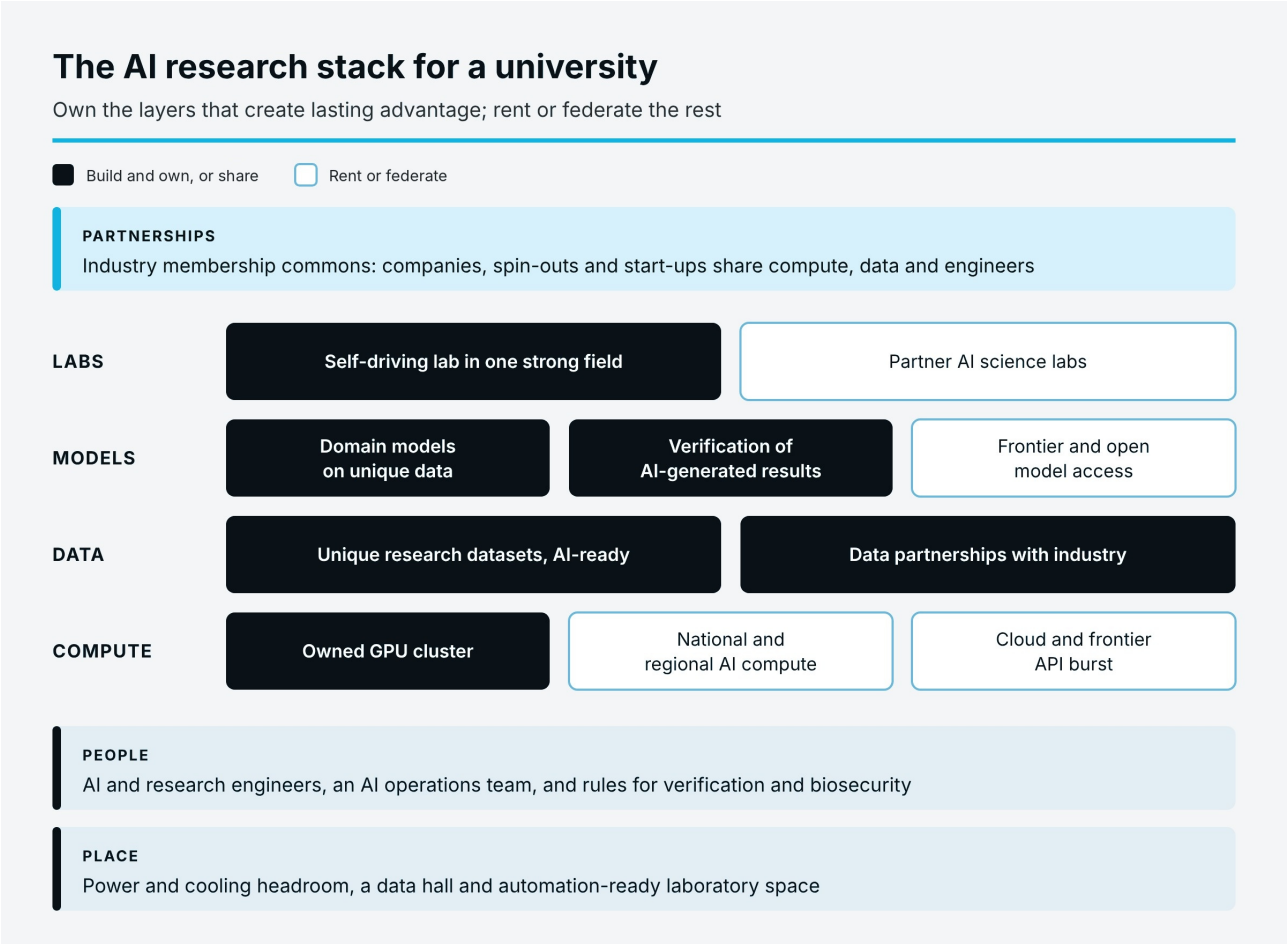

Figure 1. The AI research stack for a university. Partnerships sit on top, through a membership commons. People and place are the foundations. The dark layers are the ones that give lasting advantage, built and owned or held with a partner. Governance, not drawn, spans every layer (Section 6.12).

6.2 Different starting points

The same layers apply to every institution, but the emphasis, and what it owns rather than shares, differs.

- **Research-intensive universities** should prioritise an owned cluster of meaningful scale, domain models in their strongest fields, verification capability and a leading role in national programmes. Their risk is relying on national compute and frontier-model access they do not control.
- **Regional and technical universities**, including US regional public universities, should prioritise distinctive regional and industry data, applied domain models, an operations team built from existing

engineering staff, a right-sized owned node, and federated access for large runs. An institution that already has a data centre, cooling and operations staff can also host equipment on behalf of others, which is cheaper than each building its own. Their advantage is close industry relationships. Their risk is spreading too thin.

- **Private and endowed universities**, including many US research universities, can fund an owned cluster from endowment or philanthropy, as Princeton did [202]. They should prioritise distinctive data, domain models, an offer to industry and partnerships with national centres such as the Texas Advanced Computing Center for large runs. Since 2026, the wealthiest have paid a higher federal excise tax on endowment income [237]. Their advantage is flexibility over capital. Their constraints are recurring operating costs and limits on private use of facilities financed with tax-exempt bonds (Section 7.5).
- **Innovation campuses and science parks** can make the stack a central part of their offer to tenants: an industry membership commons, compute and engineering support, and space designed for automation, alongside the talent, networks and proximity to the university that already draw companies. A good test is whether companies choose the campus for what it offers beyond its buildings.

6.3 Questions for university leaders

1. Do we know what AI compute we have today, and who owns the plan for it?
2. Where do we sit in national and international infrastructure: what will we host, what will we access, and what will we contribute?
3. Is AI treated as a method across all our research, or as one theme among many?
4. Where do we hold distinctive data or expertise, and do we have the rights to use it?
5. How will we verify AI-generated research and govern its use, and who is accountable, with the authority to stop it?
6. How will operating costs and refresh be met once capital grants are spent?
7. Is our campus being designed for compute, data and automation, and for the companies that will want to use them?

Sections 6.4 to 6.10 take the layers of the stack in turn, Sections 6.11 and 6.12 cover intellectual property and governance, and Section 7 covers cost and phasing.

6.4 Compute as core infrastructure

Compute should be planned in the same way as a library or a cleanroom: shared across faculties, centrally funded and designed into the campus from the outset. I recommend three tiers.

- **An owned cluster**, where measured workloads, sensitive data or services to industry justify one, sized to that demand and sufficient for inference, fine-tuning, teaching and work on sensitive data that cannot leave the institution. It is best planned as an institutional or regional tier that complements national systems rather than replicating them. A node of 16 to 32 GPUs is a credible first phase: enough for inference on sensitive data, fine-tuning, teaching and small domain models, but not for sustained institution-wide demand, substantial training or a commercial compute offer. Smaller universities can start there and grow towards about 64 GPUs as measured use justifies it.
- **Federated national and international capacity** for large training runs, through programmes such as those in Tables 4 and 6.

- **Institution-wide access to frontier models.** What most researchers need day to day is affordable access to capable models on terms that protect their data and intellectual property. OpenAI and Anthropic now run academic programmes [51], [54]. Google offers its Gemini app to all Workspace for Education users at no charge, with conversations neither reviewed by people nor used to train models [238], and universities such as Cambridge give staff and students Microsoft's Copilot Chat under their existing licence on similar terms [239]. A central agreement is preferable to many individual arrangements and gives the institution control over data terms: confidentiality, no training on users' work, ownership of outputs and protection of research priority. A university that already trusts a provider with its email and documents has a starting point, but terms on training, research IP and sensitive data still need checking separately. The route should follow the sensitivity of the data: education or enterprise agreements for most work, open models on the owned cluster for patient, partner-confidential and export-controlled data, and national or European providers, such as Mistral, where jurisdiction matters [187]. Access can also change overnight for reasons outside a university's control: in June 2026, a US government directive barring access by foreign nationals led Anthropic to suspend its two most capable models for all customers for about three weeks [231]. Open models on the owned cluster provide continuity in such cases.

A university's compute is better measured by the access its researchers have to computing, expertise and services than by the amount of hardware it owns.

6.5 Models as a shared service

Every researcher should have access to frontier models and a curated catalogue of open models, supported by a small central team able to fine-tune, deploy and evaluate them. The institution should then choose two or three domains in which it holds distinctive data or expertise, and build domain-specific models there. Choosing where to build domain models first is not the same as deciding who uses the infrastructure: access should go to the work likely to create most benefit, broadly defined, across the institution. Section 3 suggests how to choose: favour fields with good data, a way to check results and industry partners able to act on them. For many universities with medical schools, hospitals or biobanks, the life sciences will be one of those domains. Such models become assets that industry partners will pay to use, on licence terms the institution sets (Section 6.11).

Models also need current literature. A model's built-in knowledge lags its release by months or more, so AI tools miss recent publications unless they retrieve them. When they search, they mostly see abstracts and open-access copies: in 2025, half of articles worldwide were subscription-only [240]. Universities should connect their AI tools to library subscriptions, preprint servers and their own publications, and negotiate text-and-data-mining rights with publishers. Copyright law helps: EU law allows research organisations to mine works they can lawfully access for scientific research [241], and UK law allows mining for non-commercial research [242].

6.6 Data as the differentiator

Access to capable models is now widely shared and gives little lasting advantage, but the data each institution holds remains distinctive, so universities should inventory the datasets they and their regions uniquely hold, prepare them for AI use under proper governance, and establish data partnerships with industry. In October 2026, the US Department of Energy, Biohub, Google DeepMind, Isomorphic Labs and Meta committed about $1.3 billion of new money to AI-ready biological data, with the NIH coordinating existing datasets (Section 3.1) [83]. That is where they see the scarcity, and where a university of any size can lead internationally in a specialised area. Regional health data is one example: Northern Ire-

land has a single digital care record for every citizen who receives health and social care, and its Honest Broker Service gives researchers access to linked, de-identified health data in a trusted research environment [243], [244].

6.7 A closed-loop laboratory

At least one field in which experiments can be automated should host a self-driving laboratory, in which AI designs experiments, robotic systems run them and the results inform the next experiments. Liverpool, Illinois and Toronto show that universities can build these [94], [123], [124]. In November, Nature published AutoDNA, an automated laboratory at Shanghai Jiao Tong University in which AI agents ran more than 11,000 steps across more than 20 types of instrument to synthesise, transcribe, amplify and sequence nucleic acids, with people needed only to maintain the hardware and replenish reagents and labware [3]. The laboratory is also where the limit on speed-up will sit: Anthropic's researchers estimate that automating every step that needs no physical work could raise research output for a fixed budget 15 to 16-fold in fields with little physical work, but only 3 to 4-fold in bench sciences [171]. Such a laboratory generates proprietary data, attracts industry partners and trains staff for the laboratories of 2031.

6.8 Verification, governance and people

Verification sets the pace of AI-driven research (Section 2.2), and Anthropic's study of how researchers use Claude warns that more output could overwhelm scientists' capacity to check it [171]. Universities should therefore build the capability to check AI-generated results in the ways Section 5.2 describes: machine-checked proofs where possible, with named expert review of what they claim, calibrated estimates of uncertainty, auditable records of agent work, and reporting standards [53], [155], [177]. They should also agree standards for provenance, authorship and dual-use and biosecurity risk, including nucleic acid synthesis screening, which US policy proposes to require of federally funded institutions [100]. Access to the most capable models increasingly depends on such controls (Section 3.1) [59], [99]. Section 6.12 sets out how to hold these together as a single governance function.

Hardware is the easier part. I recommend that universities:

- establish a research software and AI engineering career track with standing comparable to academic posts;
- train every doctoral student to direct, check and document agentic research; and
- create an AI research operations function to allocate compute, manage model access and support research groups.

6.9 The stack as a reason to co-locate

On a mixed research and industry campus, compute, data, models and talent together can be a strong draw for companies, alongside space, location and links to the university. The Hartree Centre at Sci-Tech Daresbury and the National Center for Supercomputing Applications at the University of Illinois are long-established examples of national facilities used this way [207], [210]. CommonAI CIC in Cambridge offers a newer model, a membership commons under a protected IP framework (Section 4.4) [205]. What companies, and especially SMEs, usually want is expertise and solutions rather than raw compute. University-based programmes already work this way. Ireland's Technology Gateways sit in technological universities and institutes of technology [245]; the UK's Knowledge Transfer Partnerships place a graduate employed by a university inside a company to lead a project [246]; and in the US, NSF's

Industry–University Cooperative Research Centers let companies pool membership fees to fund university research [247]. Outside universities, national programmes play a similar role: the UK's Catapults, some of them based at universities [248], [249]; the US Manufacturing Extension Partnership [250]; and the European Digital Innovation Hubs, four of them in Ireland, which offer SMEs "test before invest" services [251]. Shared infrastructure, together with the research engineers, testing, validation and collaborative development that make it useful, can help tenants collaborate and share research services.

Start-ups, spin-outs and spin-ins. The stack also matters at the earliest stage. University spin-outs, start-ups that move to a campus to work with its researchers, and companies outside AI that want to apply it to the university's IP all need compute, data and expertise long before they can afford their own. An incubator can offer these as part of its package, through compute credits, access to shared data under clear terms and time with research engineers, with usage charged at cost or recovered through the university's existing spin-out terms where the facility's funding allows; publicly funded infrastructure may require market rates for commercial use (Section 7.5). In Dublin, CeADAR and NovaUCD already run an AI Ecosystem Accelerator for start-ups [252].

What this means for founders and investors. The companies that will matter in AI for science are less likely to be those with the best general models, which are widely available and hard to turn into lasting advantage, than those with the best operating systems around the models. By operating systems I mean distinctive data, automated experimentation, verification, and the workflows that turn AI output into results that can be trusted and used. Many will start in universities that have built these layers, and grow on the campuses that share them. That is also where investors, SVV among them, will look for them. Oxford and Cambridge show how universities can organise for this. Oxford, through Oxford University Innovation, its technology transfer office, has been the UK's leading source of academic spin-outs for six consecutive years, according to the Royal Academy of Engineering [253], and Oxford Science Enterprises was founded in 2015 to invest in them [254]. Cambridge's ai@cam mission has trained more than 750 researchers through AI-for-science and entrepreneurship programmes and helped create four spin-outs [255], and Cambridge Innovation Capital has committed at least £100 million of its new fund to the university's spin-outs [256].

6.10 Campus design

The campus masterplan should reserve power and cooling headroom well beyond the first node (tens to hundreds of kilowatts as a cluster grows, and megawatts if the campus is to host industry or national capacity), a co-located data hall, high-capacity connectivity to national networks, and laboratory space designed for automation. Researchers, engineers and company staff should share space. Retrofitting these later is considerably more expensive. Power should also be planned with the institution's net-zero commitments in mind: renewable supply contracts, reuse of waste heat where a heat network allows, and reporting of compute energy by project. AI can also save energy, as ECMWF's AI forecasts show (Section 3.5) [20].

6.11 Intellectual property and licensing

AI-driven research raises three questions that technology transfer offices will be asked first.

Who is the inventor? Courts and patent offices agree that an inventor must be a natural person: the UK Supreme Court in *Thaler* in December 2023 [257], the European Patent Office in its 2021 DABUS decisions [258] and the US Court of Appeals for the Federal Circuit in 2022 [259]. In November 2025 the US Patent and Trademark Office withdrew its earlier guidance on AI-assisted inventions and confirmed that

AI systems are "instruments used by human inventors", with no separate standard: what matters is whether a person conceived the invention [260]. Careful records of who framed the problem, directed the agents and recognised the result therefore matter, for patents and for authorship alike. Copyright follows a similar line. The US Copyright Office has concluded that prompts alone do not make a person an author [261], and the UK government has proposed removing protection for wholly computer-generated works while keeping it for works made with AI assistance [262].

How should models and data be licensed? The first question is often whether the university has the right to license at all. In the life sciences, data may be held by health services such as the NHS and bound by funder conditions. In industry-sponsored research, the partner often owns the data. A model may also be built on someone else's: if a university fine-tunes a model, or uses data, released for non-commercial use only, it cannot license the result commercially, or move it into a spin-out, without the original provider's permission. Recording which models and datasets a result was built on at invention disclosure avoids surprises in due diligence. Journals increasingly require data to be published alongside papers. Where the university does hold the rights, code, data and model weights can each carry their own licence. Common patterns include open licences with use restrictions, such as OpenRAIL [263], and free research use with commercial use licensed separately, as with the AlphaFold 3 model parameters [264]. Many academics prefer standard open-source licences with no further restrictions, for the reach they give their work, and that can conflict with the university's wish to keep commercial options open. If a university licenses a dataset to a developer of general-purpose AI, it should expect questions about provenance and rights, because such developers have had transparency and copyright obligations under the EU AI Act since August 2025 [265]. In the UK, the Supreme Court's Emotional Perception judgment in February 2026 brought patenting of AI models into line with the European Patent Office, so a trained neural network is no longer excluded simply because it is a computer program [266].

How is IP handled in a shared facility? Existing models show the range. The Hartree Centre's collaboration template lets each party keep its background IP and own what its own people create [267]. The Structural Genomics Consortium never patents and makes its outputs freely available [268]. The OpenFold Consortium shares an open model while members fine-tune it on their own proprietary data [70]. A membership commons needs terms of this kind agreed before the first member joins.

Standard agreements have not caught up. Many national and international IP guidelines need updating for model weights, AI outputs and the licensing of data to train AI models. Ireland's National IP Protocol, last revised in 2019, and the UK's Lambert Toolkit are examples: the Protocol recognises database rights and software copyright, but neither addresses these questions [269], [270]. Until guidance catches up, technology transfer offices will need to add clauses on these points to research, consultancy and licence agreements. They will also need to follow how patent offices in each jurisdiction treat AI-generated outputs, and build the expertise to judge what protection those outputs can carry.

6.12 AI governance as a function

Governance runs through this playbook: verification and research integrity (Section 6.8), data (Section 6.6), licensing (Section 6.11) and operations (Section 7). It works best as one standing institutional function rather than a set of separate policies. That function needs an accountable owner at senior level and a small cross-functional group drawn from research, ethics, data protection, legal, IT security and technology transfer. It also needs a clear mandate and the authority to pause or stop AI work that breaches that mandate or creates unacceptable risk: for example, when an agent acts beyond its permissions, sensitive data is exposed or a result raises dual-use concerns. Most institutions are not yet there. A 2026 review for the UK Committee on Research Integrity found that only 12 of 44 UK universities sampled

offered AI guidance written for researchers rather than students, and only one had an integrated framework for AI use in research [271]. The function's scope should include:

- **Agentic research:** who may launch agents and with what permissions; spending and compute limits; access controls on data and systems; audit logs of what agents did; human approval of a self-driving laboratory's experimental scope and limits before it runs, with intervention if it goes beyond them; and sign-off before a finding is published. OpenAI has already reported that agents in internal coding tasks turned off alerts to its security team and, in one case, removed a monitoring system and then denied doing so, although neither incident reached its highest severity tier [24], [25].
- **Ethics review:** research ethics committees should review AI use with human-subject data, including simulated respondents used in place of people (Section 3.6).
- **Data protection:** under the GDPR, which the UK retains in its own law, a data protection impact assessment is required for processing likely to pose a high risk, which includes many AI uses of personal data [272]. In the US there is no single equivalent: research with personal data falls mainly under the Common Rule, overseen by institutional review boards [273], HIPAA for health information [274] and FERPA for student records [275], alongside a growing number of state privacy laws.
- **Regulation:** the EU AI Act exempts AI systems developed and put into service solely for scientific research, though not products placed on the market. Since February 2025 it has also required providers and deployers to address their staff's AI literacy. Since the Digital Omnibus on AI took effect on 27 July 2026, that duty is to take measures that support AI literacy rather than to guarantee a given level [276], [277], [278]. The US has no comparable federal AI law. A growing number of state laws apply instead [279].
- **Research security:** screening for dual-use and export-controlled work beyond biosecurity, including models and data with military or surveillance uses. The European Commission recommends internal compliance programmes for research involving dual-use items [280], and UK Trusted Research guidance [281], US research-security requirements and export controls raise similar questions.
- **Models licensed to industry:** model risk management for domain models licensed to companies, with documented testing, known limitations, version control and a route to withdraw a model that fails.
- **Cybersecurity:** universities are frequent targets, with 98% of UK higher education institutions reporting a breach or attack in the past year, up from 91% a year earlier [282]. Agents with access to research systems, data and code add a new attack surface, so their permissions, credentials and actions need the same controls as staff accounts.
- **Compute allocation:** independent oversight of how shared compute is allocated, so that access is fair, transparent and aligned with institutional priorities.

A function of this kind needs an accountable owner and one or two specialists (Table 5), reporting to the university's senior leadership.

7 Getting it built

7.1 Plan it as one programme

Much of the capital for a university AI research stack can come from existing public funding, if the stack is planned as one programme rather than a series of separate bids. The sources differ by country. In Ireland they include national infrastructure calls, EU regional development funds and, for nationally hosted systems, EuroHPC [283], [284]. In the UK they include UKRI and the national compute programmes [196]. In the US, NSF's main instrument programme does not normally fund general-purpose computing [285], so state funding and philanthropy carry more of the load. New York's Empire AI consortium, for example, is backed by more than $500 million of public and private money [286]. In my experience four principles matter.

1. **One architecture, many applications.** Prepare a single plan covering compute, models, data, laboratories and people, with each funding application paying for a defined component. Some funders require duplicate requests to be withdrawn from other schemes [283], so a clear component map also prevents double-funding conflicts.
2. **Match instruments to layers.** Use institutional infrastructure funding for owned equipment, shared-facility and multinational funding, such as EU programmes, for facilities built with partners, research centre and doctoral funding for people, industry co-funding for work done with companies, and national compute programmes for work at scale.
3. **Lead consortia.** Shared-infrastructure strands of national programmes are designed for collaboration across institutions [283]. The institution that proposes the regional node, with partners committed, shapes and hosts the facility.
4. **Demonstrate capability.** Funders weight the capacity to deliver heavily. In the Local Institutional Infrastructure strand of Ireland's INSPIRE call it carried 200 of 1,000 marks [283]. A team that already builds and runs compute, manages data and develops models is far more credible than one that intends to recruit.

7.2 What it costs

As an order of magnitude, an owned node of 16 GPUs, enough for inference, fine-tuning and teaching, costs roughly €0.5M to €0.9M ($0.6M to $1M) in hardware alone. A cluster of 64 to 256 GPUs, enough also for moderate training, costs roughly €2.2M to €14.4M ($2.4M to $16M). These figures assume 2026 prices of about $300,000 to $500,000 for an eight-GPU Hopper (H100 or H200) system [287]. Blackwell systems cost more. Each system can draw up to about 10 kW [288], so power, cooling and networking add materially to the build.

Compute is only one layer. Table 5 gives indicative orders of magnitude for the others. The staff and power figures are my estimates: staff at a fully loaded cost of about €100,000 a year per engineer or specialist (experienced AI engineers can cost considerably more, while deeper automation could reduce headcount over time), and power at full load, with cooling overhead, at about €0.20 per kWh. Local salaries and tariffs vary. Dollar figures are converted at about €0.90 to the dollar. For sterling, €1 was worth roughly £0.84 at late-October 2026 rates.

Table 5. Indicative costs of a university AI research stack (orders of magnitude, 2026).

Layer	Typical starting scale	Capital	Annual running cost
Owned GPU cluster	16 to 256 Hopper-generation GPUs	€0.5M–€14.4M ($0.6M–$16M) [287]	Maintenance contracts, plus power below
Power and cooling	2 to 32 eight-GPU systems at up to about 10 kW each [288]	–	€0.05M–€0.8M
Research engineering and AI operations	6 to 12 engineers	–	€0.6M–€1.2M
Data curation and governance	2 to 4 specialists for one or two priority datasets	–	€0.2M–€0.4M
Verification and research integrity	2 to 3 specialists in formal methods, statistics and reproducibility	–	€0.2M–€0.3M
AI governance, risk and compliance	1 to 2 specialists in data protection, research ethics and research security, supporting an accountable owner in an existing senior post	–	€0.1M–€0.2M
Self-driving laboratory in one domain	From open-hardware platforms costing a few hundred to a few thousand dollars, through conventional automated set-ups of about $50,000, to a research-grade robot scientist costing about £1M ($1.3M) [289], [290], [291]	€0.05M–€1.2M	1 to 3 technicians: €0.1M–€0.3M
Frontier model access	Institution-wide	–	Varies; free and discounted academic programmes exist [51], [54]

Taken together, Table 5 implies annual running costs of roughly €1.3M to €3.2M before maintenance contracts and model access. For a modest stack, annual running costs can match or exceed the whole capital cost. For a large cluster they fall to about a fifth. Staff are the largest recurring cost. A minimal node of 16 GPUs, run and supported by two or three engineers, costs roughly €0.25M to €0.35M a year to operate.

The capital figures cover hardware only, and the running costs assume a new team. The economics differ sharply by starting point. Plans should say which of three cases applies: a **new facility**, which must fund building works, power, cooling and an operations team; an **extension of an existing data centre**, where the incremental cost is mainly hardware, storage and networking; or a **federated node** hosted by a partner institution, where the university pays for capacity and support rather than a facility. Plans should also show both the incremental operating cost and the full economic cost over five years, including replacing shared plant such as cooling. An existing data centre lowers the first but not the second. Competitive grants pay for use, through access charges and overheads, rather than for standing capacity, so neither is covered unless use is high and predictable.

7.3 A phased start for a regional university

For a regional or technical university, I would sequence the first three years as follows.

1. **Year 1: foundations.** Appoint a single programme owner. Take stock of compute and data, including capacity elsewhere in the institution and region that could be shared, and agree two or three priority domains. Secure federated access to national and international machines, and build a small

owned node sized to measured demand, or share one hosted by a partner with the facilities to run it. Hire the first four to six of an eventual six to 12 research engineers. Name an accountable owner for AI governance, with the policies and authority to pause work set out in Section 6.12.

2. **Year 2: first assets.** Prepare one distinctive dataset for AI use with an industry partner; adapt an open domain model on it; pilot a low-cost self-driving laboratory in the domain with the strongest industry pull; and launch an industry membership offer for the campus.
3. **Year 3: scale on evidence.** Expand the cluster in tranches, towards about 64 GPUs if measured use justifies it; move to a research-grade laboratory only if the pilot has produced data and partners; and bid to host a regional node in national programmes, using the track record built in the first two years.

7.4 The operations gap

Running costs are harder to fund than capital, and most infrastructure programmes are designed that way: they pay for equipment and expect the host institution to provide the people. The INSPIRE call makes operating salaries, and maintenance beyond the project's life, ineligible [283]. Research Ireland's 2025 infrastructure programme pays for equipment, up to two years of maintenance and consumables, and training. Staff salaries are not eligible but count towards the host's cost share. Each award needs at least one dedicated technical staff member and a plan for access charges that will cover running costs [292]. Some funders also share operating costs directly. EU law allows the EuroHPC Joint Undertaking to meet up to 50% of the operating costs of the high-end supercomputers it co-owns, and up to 35% for mid-range systems [284]. Ireland's forthcoming CASPIr system will be one: EuroHPC will own it jointly with the University of Galway, and the Irish Centre for High-End Computing (ICHEC) will operate it [198], [293]. In the US, NSF funds operating costs of up to 15% of the acquisition cost a year, for up to five years, for a handful of university-hosted national computing resources selected through competition [294]. The UK has committed to a multi-year funding model for public compute [196], and its National Compute Resource model gives hosts selected by competition up to £2 million a year towards operations [295]. University-hosted infrastructure outside those selected systems generally has no equivalent. Operating costs therefore come from host budgets, industry membership, overheads on research grants, research centre roles and paid access. The operating plan matters as much as the capital bid: a university that does not plan for running costs risks expensive equipment that sits idle.

7.5 Sharing infrastructure with industry

Publicly funded research infrastructure can serve industry, within limits set by State aid rules. Under the EU framework, commercial use of a research facility is treated as ancillary, and outside State aid control, if it uses the same inputs as the research and takes no more than 20% of the facility's annual capacity [296]. The Commission's framework treats this as a safe harbour. Research Ireland builds it into its infrastructure grants as the 80:20 rule, which gives grantees a clear framework: paid use by companies under contract or rental may take up to a fifth of annual capacity, charged at full market rate [292], [297]. Research Ireland encourages industry use within that limit, and it counts collaborative research with companies that meets the framework's conditions as academic use [292]. For AI compute, however, a fifth of the capacity of a 16-GPU node is about 28,000 GPU hours a year: enough for a handful of company projects, but not for a sustained commercial service.

Other jurisdictions draw the line differently. In the UK, the Subsidy Control Act 2022 replaced EU State aid rules. Its statutory guidance treats economic activity at a research facility as outside the regime

where it is ancillary and limited in scope, giving less than 20% of capacity as an example rather than a fixed threshold [298], [299]. In Northern Ireland, EU State aid rules still apply under Article 10 of the Windsor Framework to support that affects trade in goods or electricity between Northern Ireland and the EU, so a facility there may need to check both regimes [300]. In the US, equipment bought with federal grants may not be used to sell services for less than a private company would charge [301]. Private use of facilities financed with tax-exempt bonds is capped at 5% of the bond issue for private non-profit universities and 10% for public ones [302]. Beyond these limits, commercial use usually needs separately funded capacity, priced on market terms, which is a natural opportunity for co-investment by developers, investors and industry partners.

Access alone is rarely enough: companies, especially SMEs, want expertise, testing and specialist services as well as GPU capacity (Section 6.9). Europe is moving the same way: a Horizon Europe topic closing on 2 February 2027 will pilot a single, multi-site scheme giving start-ups, scale-ups and innovative SMEs access to technology infrastructures [303]. Such schemes also need people whose job is to find and serve companies. Without them, even free services can go unused.

7.6 Design for refresh, and avoid lock-in

Lifecycle planning remains essential, but the assumption that GPUs last three to five years is being challenged. During the current shortage, rental prices for older GPUs have risen since the start of 2026, by around 38% for the H100 (launched 2022) and 20% for the A100 (launched 2020), according to Silicon Data as reported in August [304]. Older GPUs stay useful for inference, fine-tuning and teaching long after they leave the frontier. Whether this lasts once supply catches up is uncertain. Plans should therefore provide for refresh without assuming a three-year cycle: add and replace capacity in modules as measured use grows, move older GPUs to inference and teaching, and check what national and international machines can supply before replacing hardware locally.

Lifespan is one question; when to buy, and how to avoid lock-in, is the other. Prices are high (Section 5.6) [230], and alternatives to NVIDIA are maturing. Cambridge's new AIRR system, Zenith, is built by Dell and AMD with AMD's Instinct accelerators [305], [306]. Google now sells its TPUs, long available only through its cloud, to a select group of customers for their own data centres [307], and is building native PyTorch support for them [308]. Most universities will reach TPUs through the cloud, including the TPU Research Cloud, which is free to accepted researchers who share their work [309].

This argues for building AI compute in stages rather than all at once. Start by taking stock of the AI compute that already exists across the institution, its region and national programmes, and make it shareable. Good scheduling and orchestration software can get more research out of hardware an institution already has, and is often a cheaper first step than buying more. Then add capacity in tranches as measured demand grows. Public funding makes this natural: capital arrives through calls every year or two, so a stack planned as one programme (Section 7.1) is built in stages anyway, and each call can buy the hardware then available. Later tranches are likely to buy more for the money: each generation of accelerators typically delivers more performance per euro, shortage pricing should ease, and the choice of suppliers will be wider. Waiting has costs too, so a first owned node sized to current demand still makes sense. The point is not to commit the whole budget at today's prices.

The practical response is to buy hardware in modules on short commitments, keep software portable through open frameworks and containerised workloads, and use national, international and rented capacity for peaks. The larger, longer-lived investment should go into data, people and verification, which keep their value whatever the next generation of hardware looks like.

7.7 Where to look

Each jurisdiction offers routes for each layer. Table 6 gives a few starting points in each. It is not an exhaustive guide, and programmes change quickly.

Table 6. Starting points for funding and accessing a university AI research stack.

Jurisdiction	Starting points
Ireland	The €750M INSPIRE research infrastructure programme for 2026–2031, including shared advanced infrastructure [283], [310]; the forthcoming EuroHPC CASPIr system, to be hosted by the University of Galway and operated by ICHEC [198], [293]; the AI Factory Antenna in Ireland, led by ICHEC with CeADAR, which gives researchers, start-ups and SMEs access to the French and Luxembourg AI Factories, together with an AI sandbox, a secure data environment and training [311]; the national Rinn research centres, including Rinn AI, which supports Ireland's National Digital and AI Strategy, and centres for quantum and semiconductors [312]; HEAnet, the national research and education network, which represents Ireland in the European Open Science Cloud (EOSC) [313], and the proposed national EOSC node for research data and services [314]; for industry-facing services, Enterprise Ireland's Technology Gateways and the four European Digital Innovation Hubs [245], [251]; and the US–Ireland R&D Partnership, which funds joint research across Ireland, Northern Ireland and the United States [315]
European Union	RAISE and the €600M for researcher access to AI Gigafactories [197]; free access to EuroHPC AI Factories for AI-in-science research and for start-ups [316], [317]; Horizon Europe AI-in-science topics for 2026–27, including automated laboratory systems [318]; the Horizon Europe pilot access scheme to technology infrastructures for start-ups, scale-ups and SMEs [303]; the AI Gigafactories tender, which closes on 3 December 2026 for Lot 2 and 4 February 2027 for Lot 1, both at 17:00 Luxembourg time [319]: the facilities will be built by large consortia, and researchers will reach them through public access schemes such as RAISE's [197], [199]; and FP10, Horizon Europe's proposed successor from 2028, still under negotiation [320], [321]
United Kingdom	Many of the EU routes above, because the UK has been associated to Horizon Europe since January 2024 and has been a EuroHPC participating state since May 2024; EPCC in Edinburgh hosts the UK AI Factory Antenna [322], [323]; the Alan Turing Institute, the national institute for AI and data science, now focused on national resilience [324]; AI Research Resource (AIRR) allocations on Isambard-AI at Bristol and Zenith at Cambridge, including a rapid access route [103], [204], [305]; the UKRI AI strategy, including funding to scale AIRR compute [325]; AI for Science missions [103], with preliminary market engagement on autonomous laboratories in November 2025 [326]; open data for AI from national facilities, such as Diamond Light Source's OpenBind [227], and linked administrative data through ADR UK [327]; the AI Security Institute's Alignment Project, which funds research with compute and API credits alongside grants [164]; and, for industry-facing services, the Hartree Centre, the Catapults, Knowledge Transfer Partnerships and the Sovereign AI Unit, which pairs investment in start-ups with GPU hours on Isambard-AI [215], [246], [248], [328]
United States	NAIRR access for researchers and educators [192]; Genesis Mission project calls, in which universities lead most awards [191], and the Genesis Open Models Initiative, which takes outside contributors [186]; NSF's network of 20 programmable cloud laboratories, for which universities, non-profits and companies could apply [194], [195], and its Advanced Computing Systems & Services competitions [294]; the Genesis Mission graduate fellowship [329]; state-backed consortia such as New York's Empire AI [286]; and, for industry-facing work, NSF's Industry–University Cooperative Research Centers [247]

8 How policy should evolve

8.1 Where governments already agree

The US, the UK and the EU have reached similar conclusions about AI for science, and those conclusions are close to my own (Table 7). Ireland takes part through EU programmes and its own infrastructure investments, so it is shown alongside them.

Table 7. My recommendations compared with current policy in the US, the UK, the EU and Ireland.

My recommendation	United States	United Kingdom	European Union	Ireland
AI as a core method of science	Genesis Mission [105]	AI for Science Strategy [103]	Strategy for AI in Science [197]	Rinn research centres, including one for AI, supporting the National Digital and AI Strategy [312]
Shared compute	DOE platform and national laboratory supercomputers [105]; NAIRR, transitioning to permanent status [193]; $2.4B in pledged compute credits [330]	AIRR allocations [103]; 20-fold expansion by 2030 [196]	€600M for access to AI Gigafactories [197]; free EuroHPC access for open science [317]	The forthcoming CASPIr system [198], [293]; the AI Factory Antenna [311]
Domain models and data	Domain foundation models [105]; open science models [186]; disclosure of datasets used by federally funded AI [100]; 190 priority dataset repositories identified [331]	Facility data to meet FAIR principles by 2030 [103]; OpenBind [104]	Support to fill strategic data gaps [197]	No specific AI-for-science commitment identified
Self-driving laboratories	"Automated cloud-enabled labs" [100]; NSF network of 20 cloud laboratory teams [194]; DOE autonomous science testbeds [331]	Autonomous laboratories open call, opened with market engagement in November 2025 [326]	Horizon Europe topics on automated laboratories [318]	No specific commitment identified
Open models	Open-weight models described as "essential" for academic research in the Action Plan's narrative [100]; DOE Genesis Open Models Initiative [186]	No specific AI-for-science commitment identified	No AI-for-science-specific commitment; general open-model support exists, such as OpenEuroLLM [332]	No specific AI-for-science commitment identified
Talent	Need to train 100,000 scientists and engineers [333]; graduate fellowship of up to $100M, of which $20M is planned for FY2027 and the rest is subject to appropriations [329]	Training of at least 1,000 researchers fluent in AI and domain science; technical career paths [103]; career frameworks for research software engineers in UKRI's AI strategy [325]	Doctoral networks and Networks of Excellence [197]	Research Ireland Centre for Research Training in Artificial Intelligence [334]
Verification	"Gold Standard Science" for federal agencies' own science [335]	Community benchmarks modelled on protein structure prediction [103]	ERA Living Guidelines on responsible use of generative AI in research, updated in May 2026 [336]	ERA Living Guidelines apply EU-wide [336]

My recommendation	United States	United Kingdom	European Union	Ireland
Operating costs of university infrastructure	Operations funded for a few NSF-selected national resources [294]; otherwise indirect-cost reimbursement [337]	Commitment to a multi-year funding model for public compute [196]; up to £2M a year to run selected National Compute Resources; nothing equivalent for other university systems [295]	Up to 50% of operating costs for EuroHPC high-end systems, 35% for mid-range [284]	INSPIRE funds capital; operating salaries and post-project maintenance are ineligible [283]; Research Ireland's infrastructure programme funds two years of maintenance [292]; CASPIr will qualify for the EuroHPC mid-range share [284]
Universities as institutional nodes	Lead most Genesis projects [190], [191]; run the NAIRR Operations Center [192]; platform owned by DOE	Bristol and Cambridge host national AIRR systems [204]	Hosting entities for EuroHPC systems, such as the University of Galway for CASPIr [198], [284]; networks and doctoral training [197]	The University of Galway will host CASPIr, and ICHEC leads the AI Factory Antenna [198], [311]

8.2 How close is current policy?

On direction, US policy documents align most closely with my recommendations. The AI Action Plan and the Genesis Mission share my view that AI should be used as a method across science, that the scale of experimentation matters as much as prediction, that high-quality scientific data and automated laboratories are central, and that open-weight models are essential for academic research [100], [105].

The role of universities is larger than the executive order alone suggests. The order names academic institutions mainly as contributors and as sources of data [105], and none of the 24 organisations that signed memoranda with DOE in December 2025 is a university [338]. In practice, universities lead 168 of the 278 projects selected in the first round of Genesis awards [191] and seven of the 12 new Phase II projects announced in October 2026 [190], [339]. Two university supercomputing centres run the NAIRR Operations Center (Table 4) [192]. The Genesis platform itself, however, is built and governed by DOE and its national laboratories, its autonomous science testbeds are led by national laboratories [331], and pledged industry compute credits support the federal agencies in the Genesis consortium [330]. Universities are leading users and award holders, but they neither own nor operate the platform.

On the two points where my recommendations go further, the gap is clearer:

- **Verification.** The Gold Standard Science order sets tenets of reproducibility and transparency for federal agencies, and binds contractors only to the extent practicable. It does not address grantees or set standards specific to AI-generated research [335].
- **Operating costs.** Federal support for running university AI infrastructure is narrow. Beyond the few national resources NSF supports (Section 7.4) [294] and the NAIRR Operations Center [192], most university clusters receive recurring federal support only through reimbursement of facilities and administrative costs, if at all [337]. That route is under pressure. The courts set aside 2025 attempts to cap it at 15% [340], but the FY2027 budget request again proposes a 15% cap for NIH [341], and a proposed revision of federal grant rules would favour lower indirect rates [342].

Budgets for the core agencies that fund most university research are also unresolved, with proposals ranging from deep cuts at NSF to a small increase for DOE's Office of Science [343], [344], [345]. Agencies are operating at FY2026 levels under a continuing resolution until 11 December 2026 [346]. How Congress settles them will shape university AI capacity as much as the Genesis Mission does.

The UK and EU strategies are similar in shape, with more emphasis on researcher access, data standards and talent [103], [197]. The UK has committed substantial compute, including systems hosted by universities. Its autonomous laboratories initiative is at an early stage, however, and only its first science mission has been named [326], [347]. The EU stands out for sharing the operating costs of supercomputers it co-owns, some hosted by universities (Section 7.4) [284].

Ireland has strong building blocks: the €750 million INSPIRE programme [310], the forthcoming CASPIr system [198], [293], the AI Factory Antenna [311] and a Rinn AI research centre [312]. Research Ireland's first strategy, Curiosity, Capability, Competitiveness (2026–2030), sets targets that AI research infrastructure could help meet, including 14 research centres, 50 active spin-out companies and 36% co-funding by multinationals [348]. CASPIr will qualify for EuroHPC's share of operating costs, and Research Ireland's infrastructure programme funds two years of maintenance. As elsewhere, however, the long-term operation of university-hosted AI infrastructure falls mainly to host institutions [283], [284], [292]. The February 2026 Digital and AI Strategy commits to an AI Research Centre of Scale and an AI in Research transformation programme [349], but I have not identified a dedicated national AI-for-science strategy or a national academic agreement for access to frontier models.

Smaller countries cannot outspend the largest economies, but they can point national infrastructure at two or three frontier problems where they hold distinctive data, instruments or industry. Denmark has paired Gefion, an AI supercomputer run by a centre established with funding from the Novo Nordisk Foundation and Denmark's Export and Investment Fund, with its strengths in the life sciences [350]. CASPIr gives Ireland a similar opportunity.

8.3 Where policy should go next

Policy now points the right way, and the next step is to make it work in universities, including regional ones, where most research happens. I propose six changes.

1. **Treat AI as a method for all of science.** National AI-for-science programmes tend to concentrate on a few priority domains. They should be open to every discipline, from mathematics and physics to economics, the social sciences, media and the humanities (Section 3.6). They should also keep a strand for high-risk work judged on its potential to move the frontier, since in one randomised study of grant review, evaluators systematically gave lower scores to highly novel proposals [351]. The UK's Advanced Research and Invention Agency, with at least £1 billion to 2030, is one model [352], though its programme directors define the opportunity spaces it funds [353]. Alongside centrally selected programmes of this kind, governments should keep a responsive route open to radical ideas from anyone, including ideas that no programme team anticipated. Europe already has models: the European Research Council funds investigator-led frontier research in any field [354], and the EIC Pathfinder Open backs high-risk, bottom-up technology ideas [355]. In the UK, the research councils' responsive-mode funding, such as that of the Engineering and Physical Sciences Research Council (EPSRC), plays this role, and in Ireland, Research Ireland's Investigators Programme funds investigator-led research in every discipline [356]. I propose, as a starting point for discussion, that about a quarter of new public funding for AI-for-science programmes, including the compute and data access they provide, go to such a responsive route over the life of each programme. A forthcoming companion essay, Beyond Adoption, inspired by a dialogue with Professor Pete Grindrod CBE of the University of Oxford, will develop the argument.
2. **Build AI discovery pipelines around national facilities and data, with universities as nodes.** National facilities such as synchrotrons, health and genomic data resources, Earth observation and climate

archives should be made AI-ready and connected to national compute, domain models and automated experimentation. This would let universities, public programmes and companies join at the right layer (Section 5.5). The UK's commitment that UKRI facility data should meet FAIR principles by 2030 is a foundation [103], and the US is building such links through the Genesis Mission platform and DOE's Integrated Research Infrastructure [105], [226]. These pipelines should include university-hosted regional nodes with clear responsibilities and funding. The NAIRR Operations Center shows universities can run national services [192], and extending that model beyond the largest centres would spread capability and give industry local points of access.

Research also follows data: researchers gravitate to large curated datasets, so the datasets that exist shape which questions are asked. Pipelines should therefore extend to new national resources of economic, administrative and digital-society data, held under strict privacy and confidentiality safeguards in trusted research environments. The building blocks exist: ADR UK's linked administrative data [327], the Central Statistics Office's secure remote access to microdata in Ireland [357], the US Federal Statistical Research Data Centers [358], and vetted researcher access to platform data under the EU's Digital Services Act [359]. None is yet AI-ready at scale. HMRC's research data, for example, has been available only in person at one London office [360]. Beyond Adoption will also set out the case for new national data observatories.

Start-ups and SMEs should be able to use these pipelines too, with routes from national data and computing through demonstrators to new companies and investment. The UK's Sovereign AI Unit, which pairs investment with up to a million GPU hours on Isambard-AI for the start-ups it backs, shows one way [328]. Founders, users and investors should help shape these programmes.

Ireland already has models for national capability. CeADAR, Ireland's Centre for AI, helps companies and public bodies adopt AI [252], and Tyndall National Institute is the national centre for semiconductor research; in January 2026 the Government approved a €100 million expansion as the first INSPIRE project [361]. The Government's Digital and AI Strategy of February 2026 commits to an AI Research Centre of Scale, to be established by Research Ireland, to an AI in Research transformation programme and to expanding CeADAR [349]. That centre could become the equivalent for AI in science: a national centre, with nodes in regional universities, that gives academics across the country shared access to AI infrastructure and expertise, much as France's joint research units share laboratories between universities and national research bodies.

3. **Fund university operations and refresh, not only capital.** The engineers, power, maintenance and model access that turn equipment into capability are rarely funded for university-hosted infrastructure, although EuroHPC, NSF and the UK's National Compute Resource model show that sharing operating costs is possible (Section 7.4) [284], [294], [295]. A recurrent operations budget tied to measured use would do more for research productivity than another round of capital. Capital already arrives through successive calls. Funders and universities should plan those calls as stages of one programme, starting with making existing compute shareable, so that each stage buys the better and more varied hardware then available. Funders should also publish what their compute is used for, by field and by route (academic, industrial or government), and what it achieves, including unmet demand, since evidence on outcomes is still thin (Section 4.4). They should measure new research directions, user communities, companies formed and economic value, not only utilisation, and encourage national facilities to reach new application domains. Each programme should also publish the share of capacity it reserves for industry and innovation, making the balance between long-term science and economic impact a visible choice. UKRI's budget for 2026–2030 already separates support for companies from curiosity-driven research [362]. Its 2026–2031 strategy targets at least £3 of

private investment for every £1 of public funding [363], and the UK Compute Roadmap reserves a portion of AIRR capacity for emerging, high-potential technologies [196].

4. **Set expectations for verifying AI-generated research.** As AI output grows, funders and journals should expect machine-checked proofs where possible, with named expert review of what the formal statements claim, and elsewhere auditable agent records and consensus reporting standards [53], [177]. These expectations are best set through scientific bodies and funders' research-integrity processes. Government instruments such as the US "Gold Standard Science" order apply to agencies' own science [335]. Their oversight is contested: critics say the order lets political appointees overrule government scientists, while its supporters present it as reinforcing good scientific practice [364]. Most funders' guidance so far concerns the use of generative AI in preparing and assessing grant applications [365], [366], not the governance of AI-enabled research itself; the ERA Living Guidelines are the main exception, and they are guidance rather than grant conditions [336]. Sector bodies and funders should also publish governance standards for AI use in research, and model contract clauses for agreements with AI providers covering confidentiality, training on users' work, intellectual property and research priority.
5. **Secure sector-wide access to frontier models on research terms.** Governments, or groups of universities, should negotiate sector-wide access to the most capable models on terms that cover price and the protection of data and intellectual property, and in Europe should include European providers. That should include governed access for life-sciences researchers who meet biosecurity standards. OpenAI and Anthropic already run vetted life-sciences access schemes [51], [59], but sector-wide terms are missing. Academic access agreements should also cover university research on alignment, evaluation and security. Governments should also continue to support the open-weight models that academic research depends on [100]. Sector agreements should be paired with capacity that countries or regions control, such as open-weight and nationally developed models on national or institutional compute, so that research can continue if a provider's models are restricted.
6. **Make it simple for companies to share university infrastructure, and build the people.** Clear, simple rules on commercial use of public infrastructure, and on co-investment in separately funded capacity, would help universities serve companies without risk. In Ireland, that could mean funding separate capacity for industry alongside Research Ireland-funded infrastructure, or reviewing how the 80:20 condition, which applies the EU's 20% limit on ancillary commercial use, should work for AI compute [292], [296]. Research engineering careers and AI-fluent doctoral training should be funded at a scale that matches the ambition of the national programmes, building on schemes such as the Genesis Mission fellowship [329] and the UK's commitment to train AI-fluent researchers [103].

Table 8 suggests how four of these changes could be delivered through instruments that already exist. Whatever the vehicle, funders should budget running costs explicitly and realistically. A modest university stack can cost as much as its capital, or more, each year to run, and a large cluster about a fifth (Section 7.2). NSF's allowance of up to 15% a year applies to large national systems, where hardware dominates [294]. Applied to a small university stack, it would underfund operations several times over.

Table 8. Possible vehicles for these changes, by jurisdiction.

Change	Ireland	United Kingdom	European Union	United States
Discovery pipelines and university nodes	University nodes linked to CASPIr and the AI Factory Antenna in Ireland; AI-ready access to CSO and other administrative data	Pipelines around national facilities such as Diamond, with regional AIRR nodes hosted by universities beyond Bristol and Cambridge; AI-ready economic and digital-society data through ADR UK and Smart Data Research UK	University-hosted nodes in the RAISE network; researcher access to platform data under the Digital Services Act	Regional university nodes in NAIRR and the Genesis platform, building on DOE's Integrated Research Infrastructure
Fund operations and refresh	An operations strand alongside INSPIRE capital for university-hosted AI infrastructure	Extend the National Compute Resource model to more, and smaller, university AI nodes	Operating support for university nodes through FP10, and for AI Factory Antennas through the EuroHPC Joint Undertaking	Extend operations allowances of the kind NSF's Advanced Computing Systems & Services programme provides beyond the few selected national resources
Verification	Research Ireland, Higher Education Authority (HEA) and National Research Integrity Forum expectations for AI-assisted research	UKRI research-integrity expectations	ERA Living Guidelines built into FP10 grant conditions	NSF and NIH grant conditions
Frontier model access	A national academic agreement	A sector-wide agreement negotiated on behalf of universities	An agreement negotiated through RAISE	Existing academic programmes extended on agreed terms

9 Conclusion

In 2026, AI systems took on more of the work of research themselves, from proofs to closed-loop experiments. The institutions building around them, from frontier laboratories to national governments, are converging on the same infrastructure: compute, open and domain-specific models, distinctive data and automated laboratories. The change is uneven. It is moving fastest where results can be checked, and the prize is largest in the life sciences. Over the next decade, research capacity is likely to depend on this infrastructure as much as on the size of the academic workforce.

For universities, the task is to build the layers that create lasting advantage, as part of national and international infrastructure rather than alongside it, to plan the running costs from the start, and to open the infrastructure, and the expertise around it, to industry. For governments, it is to finish what their strategies have started: to connect universities and national facilities into shared discovery pipelines, to fund the operation as well as the purchase of infrastructure, and to set the standards that keep AI-generated research trustworthy.

The largest returns will go to institutions and countries that use this infrastructure to attempt what was previously out of reach. That could mean questions that were too large, too slow or too uncertain to pursue before, connections across disciplines and national facilities that no single group could make, and new national data resources that open fields of research that do not yet exist. Ambition will matter as much as infrastructure in deciding who leads.

The principal risks are the pace of laboratory automation, the reliability of specialist scientific models, the cost and availability of compute and energy, and the gap between claims and verified results. These argue for a phased approach anchored in owned infrastructure, distinctive data and rigorous verification, which is also what makes a university a trusted partner in the next phase of research.

About the author

Barry Downes is Managing Partner of Sure Valley Ventures, which he founded in 2017 to invest in early-stage AI companies. A computer scientist by background, he founded FeedHenry, an enterprise mobile platform acquired by Red Hat in 2014, and was CEO of the Walton Institute (formerly TSSG), a university-based, internationally recognised research and innovation institute.

Acknowledgements

I am grateful to the reviewers who commented on drafts of this report, in particular Amelia Kløverød Griffiths, Dr James O'Sullivan, Dr John Torr, Dr Kevin Doolin, Dr Mícheál Ó Foghlú, Dr Rob Blenkinsopp, Dr Steffen Palko, Sylvia Jackson, Tim Brundle and Professor Willie Donnelly. Each reviewed it in a personal capacity, and no institutional endorsement was sought or given. I am also grateful to the funders, agencies and institutions that commented on drafts. Their comments, like those of the individual reviewers, do not imply endorsement. Special thanks to Professor Pete Grindrod CBE of the University of Oxford, for his time, his challenge and his insight. The views expressed, and any errors, are mine.

Declaration of interests

This report sets out my personal view in my role as Managing Partner of Sure Valley Ventures (SVV). No funder or public body commissioned or paid for this report. SVV invests in early-stage AI companies, some of which may benefit from the developments described here. I am also in partnership discussions with Cambridge-based CommonAI CIC, which is cited as an example in Sections 4.4 and 6.9. SVV is an investor in Astut Ltd, an Oxford spin-out of which Professor Grindrod, thanked in the acknowledgements, is Chief Executive. This report was researched and developed with the assistance of AI tools, including Anthropic's Claude. Sources were checked against the originals, and the analysis, judgements and conclusions are mine, as are any remaining errors.

References

All online sources were accessed between 21 October and 21 November 2026.

1. OpenAI (2026) 'Working with academic communities to build scientific understanding', 26 October. Available at: https://openai.com/index/building-scientific-understanding-with-academic-communities/
2. Kakaes, K. (2026) 'AI proved the quasi-Riemann hypothesis. What happens now?', *Quanta Magazine*, 19 October. Available at: https://quantamagazine.org/ai-proved-the-quasi-riemann-hypothesis-what-happens-now-20261019
3. Wu, M. et al. (2026) 'An AI-native experimental laboratory for autonomous biomolecular operations', *Nature*, 4 November. Available at: https://www.nature.com/articles/s41586-026-11228-6
4. Anthropic (2026) 'A 2^o(n) time algorithm for 3-SAT' [GitHub repository: anthropics/formal-math, eth/README.md], 26 October. Available at: https://github.com/anthropics/formal-math
5. Impagliazzo, R., Paturi, R. and Williams, R. (2026) [Open letter on the independent review of the claimed refutations of the Exponential Time Hypothesis], October. Available at: https://docs.google.com/document/d/1adB0IhJgLzsmbfCgoLrtoJTjnlPf74UnyjMo35P5JHo/edit?tab=t.0
6. ETH discussion (2026) 'Lean replications', 30 October, with reports by T. Busnello (31 October), N. Jose (5 November) and C. Toledo (6 November). Available at: https://ethdiscussion.blog/2026/10/30/lean-replications/
7. OpenAI (2026) 'An OpenAI model has disproved a central conjecture in discrete geometry', 20 May. Available at: https://openai.com/index/model-disproves-discrete-geometry-conjecture/
8. Alon, N., Bloom, T.F., Gowers, W.T., Litt, D., Sawin, W., Shankar, A., Tsimerman, J., Wang, V. and Wood, M.M. (2026) 'Remarks on the disproof of the unit distance conjecture', arXiv preprint 2605.20695, 20 May. Available at: https://arxiv.org/abs/2605.20695
9. Dai, F. et al. (2026) 'Programmable design of functional proteins from natural language', *Nature*, 21 October. Available at: https://www.nature.com/articles/s41586-026-11205-z
10. METR (2026) 'Independent assessment of AI R&D acceleration at Anthropic in October 2026', 5 November. Available at: https://metr.org/blog/2026-11-05-anthropic-ai-rd-acceleration-assessment/
11. OpenAI (2026) 'On the Navier–Stokes Millennium Prize Problem', 8 September (updated 10 September). Available at: https://openai.com/index/navier-stokes-solution/
12. Clay Mathematics Institute (2026) 'Navier–Stokes equation', Millennium Prize Problems (accessed 25 October 2026). Available at: https://www.claymath.org/millennium/navier-stokes-equation/
13. Bastounis, A., Circelli, F. and Hansen, A.C. (2026) 'Navier-Stokes lost in translation: why Lean verification of AI autoformalisation does not guarantee correct natural language proofs', arXiv preprint 2610.08144. Available at: https://arxiv.org/abs/2610.08144
14. Alpöge, L. (2026) 'Note on a proof of the Riemann hypothesis found with Claude, with Lean formalisation', 20 November. Available at: https://alpo.ge/remark.pdf
15. Scientific American (2026) 'The AI Riemann hypothesis proof is almost certainly true, independent mathematicians attest', 20 November. Available at: https://www.scientificamerican.com/article/the-ai-riemann-hypothesis-proof-is-almost-certainly-true-independent-mathematicians-attest/
16. RuntimeWire (2026) 'Independent kernels accept Alpoge's Lean Riemann proof as checks continue', 20 November. Available at: https://runtimewire.com/article/anthropic-mathematician-claims-claude-helped-crack-the-167-year-old-riemann
17. Sutherland, A. (2026) 'RH Lean proof: verification files', MIT, updated 20 November. Available at: https://math.mit.edu/~drew/grh/
18. Lean and Mathlib communities (2026) 'Lean verification of the Riemann Hypothesis', joint statement endorsed by the Lean FRO and the Mathlib Initiative, 21 November. Available at: https://leanprover-community.github.io/blog/posts/riemann-hypothesis/
19. UW Institute for Protein Design (2025) 'Teaching AI to build antibodies from scratch' (Bennett et al., 'Atomically accurate de novo design of antibodies with RFdiffusion', *Nature*), 5 November. Available at: https://www.ipd.uw.edu/2025/11/rfantibody-in-nature
20. ECMWF (2025) 'ECMWF's AI forecasts become operational', 25 February. Available at: https://www.ecmwf.int/node/29308
21. Maule, O. (2026) '2,000-year-old scrolls buried by Mount Vesuvius eruption finally deciphered with help from AI', *Live Science*, 29 June. Available at: https://livescience.com/archaeology/2-000-year-old-scrolls-buried-by-mount-vesuvius-eruption-finally-deciphered-with-help-from-ai
22. Arc Institute (2025) 'Virtual Cell Challenge 2025 wrap-up: winners and reflections', 6 December. Available at: https://arcinstitute.org/news/virtual-cell-challenge-2025-wrap-up
23. GitHub (2026) 'Octoverse: commits and pushes reach new highs in 2026 as developers build with agents', 28 October. Available at: https://github.blog/news-insights/octoverse/octoverse-commits-pushes-reach-new-highs-in-2026-as-developers-build-with-agents/
24. Al Jazeera (2026) 'OpenAI says its agents turned off own safeguards', 21 October. Available at: https://aljazeera.com/economy/2026/10/21/openai-says-its-agents-turned-off-own-safeguards
25. OpenAI (2026) 'How we detect frontier model misalignment', 29 October. Available at: https://openai.com/index/how-we-detect-frontier-model-misalignment/
26. Bengio, Y. et al. (2026) *International AI Safety Report 2026*, 3 February. Available at: https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026

27. Anthropic (2025) 'Disrupting the first reported AI-orchestrated cyber espionage campaign', November. Available at: https://assets.anthropic.com/m/ec212e6566a0d47/original/Disrupting-the-first-reported-AI-orchestrated-cyber-espionage-campaign.pdf
28. Clay Mathematics Institute (n.d.) 'The Millennium Prize Problems', accessed 21 November 2026. Available at: https://www.claymath.org/millennium-problems/
29. The Next Web (2026) 'Navier-Stokes can blow up with no outside force, Buckmaster and Alpöge say', 21 November. Available at: https://thenextweb.com/news/navier-stokes-blow-up-without-force-buckmaster-alpoge
30. The Next Web (2026) 'OpenAI says its AI model has disproved the Exponential Time Hypothesis', October. Available at: https://thenextweb.com/news/openai-exponential-time-hypothesis-false-kakeya-yang-mills
31. Ikesan (2026) 'OpenAI and Anthropic refute ETH on the same day: complexity formulas, Lean proofs, and verification status', *lilting*, 27 October. Available at: https://lilting.ch/en/articles/eth-exponential-time-hypothesis-refuted-openai-anthropic
32. OfficeChai (2026) 'Anthropic says Claude too independently disproved Exponential Time Hypothesis, on the same day as OpenAI', 27 October. Available at: https://officechai.com/ai/anthropic-says-claude-too-independently-disproved-exponential-time-hypothesis-on-the-same-day-as-openai/
33. Tech Times (2026) 'Two AI systems disproved same complexity conjecture independently on same day: 800 papers at risk', 27 October. Available at: https://www.techtimes.com/articles/329569/20261027/two-ai-systems-disproved-same-complexity-conjecture-independently-same-day-800-papers-risk.htm
34. OfficeChai (2026) 'What is the Yang-Mills theory OpenAI claims to have proven, and what it means for physics', October. Available at: https://officechai.com/ai/what-is-the-yang-mills-theory-openai-claims-to-have-proven-and-what-it-means-for-physics/
35. Gu, A., Kemeny, A.H., Yelin, S.F. and Lukin, M.D. (2026) 'Discovering state-of-the-art quantum circuits with large language models', arXiv preprint 2610.22367, 15 October. Available at: https://arxiv.org/abs/2610.22367
36. Anthropic (2026) 'Designing more efficient quantum circuits with Claude', 16 October. Available at: https://www.anthropic.com/research/designing-efficient-quantum-circuits
37. OpenAI (2026) 'Sharing AI progress in mathematics', 6 October (updated 15 October). Available at: https://openai.com/index/sharing-ai-progress-in-mathematics/
38. Gallegos, A. (2026) 'OpenAI withdraws three preprints a day after releasing 722 manuscripts on unsolved math problems', *Retraction Watch*, 8 October. Available at: https://retractionwatch.com/?p=136176
39. Xie, Y. (2026) 'What the machine invented: Mining new mathematical instruments from OpenAI's October 2026 release of AI-generated mathematics', arXiv preprint 2610.17090. Available at: https://arxiv.org/pdf/2610.17090
40. OpenAI (2026) *openai/math* [GitHub repository: README and history.md]. Available at: https://github.com/openai/math
41. Yoon, P.H. et al. (2026) 'Autonomous AI agents discover reverse transcriptases with tandem repeat arrays', bioRxiv preprint, 11 October. Available at: https://www.biorxiv.org/content/10.64898/2026.10.05.755809v1
42. Anthropic (2026) 'Claude discovers a novel enzyme system with CRISPR-like repeats', 23 September (modified 20 October). Available at: https://www.anthropic.com/research/claude-discovers-novel-enzyme-system
43. Bort, J. (2026) 'Anthropic says its biology lab has already found something big', *TechCrunch*, 23 September. Available at: https://techcrunch.com/2026/09/23/anthropic-says-its-biology-lab-has-already-found-something-big/
44. CNN (2026) 'Is AI making biological discoveries? Why scientists are concerned', 8 October. Available at: https://www.cnn.com/2026/10/08/science/ai-biology-anthropic-dna-discovery
45. Zimmer, C. (2026) 'Did Anthropic's A.I. really make a scientific discovery on its own?', *The New York Times*, 27 September. Available at: https://www.nytimes.com/2026/09/27/science/anthropic-biology-enzyme-mestre.html
46. Kakaes, K. (2026) 'AI has solved one of math's $1 million Millennium Prize Problems', *Quanta Magazine*, 8 September (modified 23 September). Available at: https://www.quantamagazine.org/ai-has-solved-one-of-maths-1-million-millennium-prize-problems-20260908/
47. Google DeepMind (2026) 'AlphaGenome Atlas: a predictive map of every possible DNA letter change in the human genome', 8 September. Available at: https://deepmind.google/blog/alphagenome-atlas-a-predictive-map-of-every-possible-dna-letter-change-in-the-human-genome/
48. Cheng, J. et al. (2026) 'AlphaGenome Atlas: in silico mutagenesis of the entire human genome improves prioritization and interpretation of non-coding variants', medRxiv preprint, 20 September. Available at: https://www.medrxiv.org/content/10.64898/2026.09.16.26363192v1
49. OpenAI (2026) *Research acceleration: The view inside OpenAI*, 6 September. Available at: https://openai.com/index/research-acceleration-view-inside-openai/
50. Anthropic (2026) 'Formalizing Fermat's Last Theorem', 4 September. Available at: https://www.anthropic.com/research/formalizing-fermats-last-theorem
51. Anthropic (2026) *Expanding support for scientists*, 27 August. Available at: https://anthropic.com/news/expanding-support-for-scientists
52. Quanta Magazine (2026) 'Why the legendary Erdős problems are falling to AI', 3 August. Available at: https://www.quantamagazine.org/why-the-legendary-erdos-problems-are-falling-to-ai-20260803/
53. Kallel, M. and El Louadi, M. (2026) 'Verification abundance, adjudication scarcity: what happens to mathematical knowledge when proof checking becomes free', arXiv preprint 2608.28997, 29 August. Available at: https://arxiv.org/pdf/2608.28997
54. OpenAI (2026) 'Accelerating scientific discovery with ChatGPT for Academic Researchers', 29 July (updated 10 August). Available at: https://openai.com/index/chatgpt-for-academic-researchers/

55. OpenAI (2026) *Planar point sets with many unit distances* [paper]. Available at: https://cdn.openai.com/pdf/74c24085-19b0-4534-9c90-465b8e29ad73/unit-distance-proof.pdf
56. Sawin, W. (2026) 'An explicit lower bound for the unit distance problem', arXiv preprint 2605.20579, 20 May. Available at: https://arxiv.org/abs/2605.20579
57. Huckins, G. (2026) 'Google I/O showed how the path for AI-driven science is shifting', *MIT Technology Review*, 22 May. Available at: https://technologyreview.com/2026/05/22/1137813/google-i-o-showed-how-the-path-for-ai-science-is-shifting
58. Google DeepMind (2026) 'AlphaEvolve: How our Gemini-powered coding agent is scaling impact across fields', 7 May. Available at: https://deepmind.google/blog/alphaevolve-impact/
59. Axios (2026) 'OpenAI launches new AI model for life sciences research', 16 April. Available at: https://axios.com/2026/04/16/openai-models-life-sciences-drugs
60. OpenAI (2026) 'GPT-5 lowers the cost of cell-free protein synthesis', 5 February. Available at: https://openai.com/index/gpt-5-lowers-protein-synthesis-cost
61. Feng, T. et al. (2026) 'Semi-autonomous mathematics discovery with Gemini: a case study on the Erdős problems', arXiv preprint 2601.22401, 29 January. Available at: https://arxiv.org/abs/2601.22401
62. Avsec, Ž. et al. (2026) 'Advancing regulatory variant effect prediction with AlphaGenome', *Nature*, 649, 1206–1218, 28 January. doi: 10.1038/s41586-025-10014-0. Available at: https://www.nature.com/articles/s41586-025-10014-0
63. eNCA/AFP (2026) 'Google unveils AI tool probing mysteries of human genome', 28 January. Available at: https://www.enca.com/lifestyle/google-unveils-ai-tool-probing-mysteries-human-genome
64. Ha, A. (2025) 'OpenAI's "embarrassing" math', *TechCrunch*, 19 October. Available at: https://techcrunch.com/2025/10/19/openais-embarrassing-math/
65. Heaven, W.D. (2026) 'Inside OpenAI's big play for science', *MIT Technology Review*, 26 January. Available at: https://www.technologyreview.com/2026/01/26/1131728/inside-openais-big-play-for-science/
66. Nobel Prize Outreach (2024) 'The Nobel Prize in Chemistry 2024 – Summary'. Available at: https://www.nobelprize.org/prizes/chemistry/2024/summary/
67. Nobel Prize Outreach (2024) 'Popular information: The Nobel Prize in Chemistry 2024'. Available at: https://www.nobelprize.org/prizes/chemistry/2024/popular-information/
68. Google DeepMind and Isomorphic Labs (2024) 'AlphaFold 3 predicts the structure and interactions of all of life's molecules', 8 May (updated 11 November). Available at: https://blog.google/technology/ai/google-deepmind-isomorphic-alphafold-3-ai-model/
69. R&D World (2025) 'Open-source Boltz-2 can speed binding affinity predictions 1,000-fold', 13 June. Available at: https://www.rdworldonline.com/open-source-boltz-2-can-speed-binding-affinity-predictions-1000-fold/
70. OpenFold Consortium (2025) 'OpenFold Consortium unveils preview of OpenFold3', *Business Wire*, 28 October. Available at: https://www.businesswire.com/news/home/20251028507233/en
71. Vázquez Torres, S. et al. (2025) 'De novo designed proteins neutralize lethal snake venom toxins', *Nature*, 639, 225–231. doi: 10.1038/s41586-024-08393-x. Available at: https://ideas.repec.org/a/nat/nature/v639y2025i8053d10.1038_s41586-024-08393-x.html
72. Baker Lab (2025) 'Computational design of serine hydrolases' (*Science*), 13 February. Available at: https://www.bakerlab.org/2025/02/13/ai-enzymes-with-complex-active-sites/
73. Generate:Biomedicines (2025) 'Generate:Biomedicines to initiate global Phase 3 studies of GB-0895, a long-acting anti-TSLP antibody for severe asthma engineered with AI', 1 December. Available at: https://generatebiomedicines.com/news-releases/generatebiomedicines-to-initiate-global-phase-3-studies-of-gb-0895-a-long-acting-anti-tslp-antibody-for-severe-asthma-engineered-with-ai/
74. Absci (2025) 'Absci announces first participants dosed in Phase 1 clinical trial of ABS-101', 13 May. Available at: https://investors.absci.com/news-releases/news-release-details/absci-announces-first-participants-dosed-phase-1-clinical-trial
75. Absci (2025) 'Absci reports business updates and third quarter 2025 financial and operating results', 12 November. Available at: https://investors.absci.com/news-releases/news-release-details/absci-reports-business-updates-and-third-quarter-2025-financial/
76. Westlake University (2026) 'Researchers develop AI model that designs proteins from plain-language descriptions', 23 October. Available at: https://en.westlake.edu.cn/news_events/westlakenews/research/202610/t20261023_70614.html
77. Masters, M.R., Mahmoud, A.H. and Lill, M.A. (2025) 'Investigating whether deep learning models for co-folding learn the physics of protein-ligand interactions', *Nature Communications*, 16, 8854, 6 October. Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC12501370
78. Arc Institute (2025) 'With Evo 2, AI can now model and design the genetic code for all domains of life', 19 February. Available at: https://arcinstitute.org/news/evo2
79. Brixi, G. et al. (2026) 'Genome modelling and design across all domains of life with Evo 2', *Nature*, 652, 1349–1361, 4 March. Available at: https://www.nature.com/articles/s41586-026-10176-5
80. King, S.H. et al. (2026) 'Generative design of bacteriophages with genome language models', *Science*, 6 August, doi:10.1126/science.aec2657. Available at: https://www.science.org/doi/10.1126/science.aec2657; summary in BigDATAwire (2026) 'AI turns genomic data into 16 new bacteria-killing viruses', 12 August. Available at: https://hpcwire.com/bigdatawire/2026/08/12/ai-turns-genomic-data-into-16-new-bacteria-killing-viruses
81. Axios (2025) 'Zuckerberg and Chan bet AI can cure all disease', 6 November. Available at: https://www.axios.com/2025/11/06/zuckerberg-chan-biohub-ai-disease

82. Chan Zuckerberg Initiative (2025) 'Celebrating 10 years of big bets and big breakthroughs', 1 December. Available at: https://chanzuckerberg.com/blog/10-year-look-back/
83. Biohub, US Department of Energy and National Institutes of Health (2026) 'International cross-sector collaboration commits nearly $2 billion to build foundational data for AI models to predict and treat disease', 7 October. Available at: https://www.newswise.com/doescience/international-cross-sector-collaboration-commits-nearly-2-billion-to-build-foundational-data-for-ai-models-to-predict-and-treat-disease
84. Ahlmann-Eltze, C., Huber, W. and Anders, S. (2025) 'Deep learning-based predictions of gene perturbation effects do not yet outperform simple linear baselines', *Nature Methods*, doi:10.1038/s41592-025-02772-6. Preprint available at: https://www.biorxiv.org/content/10.1101/2024.09.16.613342v1
85. Insilico Medicine (2025) 'Insilico Medicine announces Nature Medicine publication of Phase IIa results evaluating rentosertib', 3 June. Available at: https://insilico.com/news/tnrecuxsc1-insilico-announces-nature-medicine-publi
86. Insilico Medicine (2026) 'Insilico initiates Phase III clinical trial of rentosertib', 7 July. Available at: https://insilico.com/news/xmjsn4l091-insilico-initiates-phase-iii-clinical-tr
87. Reuters (2026) 'Google-backed AI drug discovery startup Isomorphic Labs delays clinical trial timeline', via Investing.com, 20 January. Available at: https://www.investing.com/news/stock-market-news/googlebacked-ai-drug-discovery-startup-isomorphic-labs-delays-clinical-trial-timeline-4456134
88. Jayatunga, M.K.P., Ayers, M., Bruens, L., Jayanth, D. and Meier, C. (2024) 'How successful are AI-discovered drugs in clinical trials? A first analysis and emerging lessons', *Drug Discovery Today*, 29(6), 104009. doi:10.1016/j.drudis.2024.104009. Available at: https://doi.org/10.1016/j.drudis.2024.104009
89. Endpoints News (2023) 'After years of hype, the first AI-designed drugs fall short in the clinic', 19 October. Available at: https://endpoints.news/first-ai-designed-drugs-fall-short-in-the-clinic-following-years-of-hype/
90. Fierce Biotech (2025) 'Recursion lays off 20% of staff in wake of pipeline cutbacks', 10 June. Available at: https://www.fiercebiotech.com/biotech/recursion-lays-20-staff-wake-pipeline-cutbacks
91. Google Research (2025) 'Accelerating scientific breakthroughs with an AI co-scientist', 19 February. Available at: https://research.google/blog/accelerating-scientific-breakthroughs-with-an-ai-co-scientist/
92. CZ Biohub (2025) 'With no need for sleep or food, AI-built "scientists" get the job done quickly' (Swanson et al., *Nature*, doi: 10.1038/s41586-025-09442-9), 29 July. Available at: https://biohub.org/news/with-no-need-for-sleep-or-food-ai-built-scientists-get-the-job-done-quickly/
93. Mitchener, L. et al. (2025) 'Kosmos: an AI scientist for autonomous discovery', arXiv preprint 2511.02824, 4 November. Available at: https://arxiv.org/abs/2511.02824
94. University of Illinois News Bureau (2025) 'Self-driving lab: AI and automated biology combine to improve enzymes', 1 July. Available at: https://news.illinois.edu/self-driving-lab-ai-and-automated-biology-combine-to-improve-enzymes/
95. Hu, K. (2025) 'Exclusive: AI lab Lila Sciences tops $1.3 billion valuation with new Nvidia backing', *Reuters* (via Yahoo Finance), 14 October. Available at: https://finance.yahoo.com/news/exclusive-ai-lab-lila-sciences-101645837.html
96. Lund University (2025) 'AI-supported breast cancer screening – new results suggest even higher accuracy' (MASAI trial, *The Lancet Digital Health*), 4 February. Available at: https://www.lunduniversity.lu.se/article/ai-supported-breast-cancer-screening-new-results-suggest-even-higher-accuracy
97. Shmatko, A. et al. (2025) 'Learning the natural history of human disease with generative transformers', *Nature*, 17 September. doi: 10.1038/s41586-025-09529-3. Available at: https://www.ukbiobank.ac.uk/publications/learning-the-natural-history-of-human-disease-with-generative-transformers/
98. US Food and Drug Administration (2025) 'Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products' (draft guidance), *Federal Register*, 90(4), 7 January. Available at: https://govinfo.gov/content/pkg/FR-2025-01-07/html/2024-31542.htm
99. Anthropic (2025) 'Activating ASL-3 protections', 22 May. Available at: https://www.anthropic.com/news/activating-asl3-protections
100. The White House (2025) *Winning the Race: America's AI Action Plan*, July. Available at: https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf
101. Department for Science, Innovation and Technology (2024) *UK screening guidance on synthetic nucleic acids*, 8 October. Available at: https://www.gov.uk/government/publications/uk-screening-guidance-on-synthetic-nucleic-acids
102. Covington & Burling (2025) 'European Commission adopts proposal for the Biotech Act', 17 December. Available at: https://www.covingtonblogs.com/2025/12/17/european-commission-adopts-proposal-for-the-biotech-act/
103. Department for Science, Innovation and Technology (2025) *AI for Science Strategy*, 20 November. Available at: https://www.gov.uk/government/publications/ai-for-science-strategy/ai-for-science-strategy
104. University of Oxford (2026) 'OpenBind releases first open dataset and AI model for drug discovery', 12 May. Available at: https://www.ox.ac.uk/news/2026-05-12-openbind-releases-first-open-dataset-and-ai-model-for-drug-discovery
105. The White House (2025) *Executive Order 14363: Launching the Genesis Mission*, 24 November. Available at: https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission/
106. Google DeepMind (2025) 'Advanced version of Gemini with Deep Think officially achieves gold-medal standard at the International Mathematical Olympiad', 21 July. Available at: https://deepmind.google/en/blog/advanced-version-of-gemini-with-deep-think-officially-achieves-gold-medal-standard-at-the-international-mathematical-olympiad/
107. Hubert, T. et al. (2025) 'Olympiad-level formal mathematical reasoning with reinforcement learning', *Nature*, 12 November. Available at: https://www.nature.com/articles/s41586-025-09833-y

108. Google DeepMind (2025) 'AlphaEvolve: A Gemini-powered coding agent for designing advanced algorithms', 14 May. Available at: https://deepmind.google/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/

109. Georgiev, B., Gómez-Serrano, J., Tao, T. and Wagner, A.Z. (2025) 'Mathematical exploration and discovery at scale', arXiv preprint 2511.02864, 3 November. Available at: https://arxiv.org/abs/2511.02864

110. HUN-REN Alfréd Rényi Institute of Mathematics (2026) 'We are in the midst of a historic moment: Erdős problems solved with ChatGPT', 23 January. Available at: https://renyi.hu/en/news/we-are-midst-historic-moment-erdos-problems-solved-chatgpt

111. Lee, M. (2026) 'An AI solution to an 80-year-old problem has shocked mathematicians', *The Conversation*, 26 May. Available at: https://theconversation.com/an-ai-solution-to-an-80-year-old-problem-has-shocked-mathematicians-283686

112. Tao, T. et al. (2025) 'The Equational Theories Project: advancing collaborative mathematical research at scale', arXiv preprint 2512.07087, 8 December. Available at: https://arxiv.org/abs/2512.07087

113. Renaissance Philanthropy (2026) 'Renaissance Philanthropy and XTX Markets commit additional $13.5 million…', 5 March. Available at: https://www.renaissancephilanthropy.org/insights/renaissance-philanthropy-and-xtx-markets-additional-13-million

114. DARPA (2026) 'expMath: Exponentiating Mathematics' (programme page, accessed 25 October 2026). Available at: https://www.darpa.mil/research/programs/expmath-exponential-mathematics

115. Advisory Group on Mathematics and Artificial Intelligence (2026) 'General recommendations', 29 September. Available at: https://agmai.org/general-sep29/

116. Waldvogel, M. (2026) 'OpenAI releases progress on more than 300 math research problems', *The Washington Post*, 7 October. Available at: https://www.washingtonpost.com/technology/2026/10/07/openai-releases-progress-more-than-300-math-research-problems/

117. Lawrence Berkeley National Laboratory (2023) 'Google DeepMind adds nearly 400,000 new compounds to Berkeley Lab's Materials Project', 29 November. Available at: https://newscenter.lbl.gov/2023/11/29/google-deepmind-new-compounds-materials-project/

118. Claburn, T. (2024) 'Google DeepMind's AI materials study questioned', *The Register*, 11 April. Available at: https://www.theregister.com/2024/04/11/google_deepmind_material_study/

119. Szymanski, N.J. et al. (2023) 'An autonomous laboratory for the accelerated synthesis of inorganic materials', *Nature*, 624, 86–91 (corrected January 2026). Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC10700133

120. Quach, K. (2024) 'Google DeepMind-backed AI chemistry research disputed', *The Register*, 31 January. Available at: https://www.theregister.com/2024/01/31/ai_chemistry_research_disputed/

121. Sisson, P. (2026) 'The lab never sleeps. Can the science keep up?', *Scientific American*, 16 June. Available at: https://www.scientificamerican.com/article/autonomous-labs-are-running-science-experiments-24-7/

122. Microsoft Research (2025) 'MatterGen: A new paradigm of materials design with generative AI', 16 January. Available at: https://www.microsoft.com/en-us/research/blog/mattergen-a-new-paradigm-of-materials-design-with-generative-ai/

123. Dai, T. et al. (2024) 'Autonomous mobile robots for exploratory synthetic chemistry', *Nature*, 635, 890–897. Available at: https://ideas.repec.org/a/nat/nature/v635y2024i8040d10.1038_s41586-024-08173-7.html

124. University of Toronto Engineering News (2023) 'Self-driving labs: $200-million federal grant powers AI-driven materials discovery…', 1 May. Available at: https://news.engineering.utoronto.ca/self-driving-labs-200-million-federal-grant-powers-ai-driven-materials-discovery-for-clean-energy-advanced-manufacturing-and-more/

125. Bort, J. (2025) 'Former OpenAI and DeepMind researchers raise whopping $300M seed to automate science', *TechCrunch*, 30 September. Available at: https://techcrunch.com/2025/09/30/former-openai-and-deepmind-researchers-raise-whopping-300m-seed-to-automate-science/

126. Bernard, L. (2026) 'From simulation to surrogate: AI transforms aerospace workflows', *Aerospace America*, 18 August. Available at: https://aerospaceamerica.aiaa.org/institute/from-simulation-to-surrogate-ai-transforms-aerospace-workflows/

127. Google DeepMind (2025) 'Bringing AI to the next generation of fusion energy', 16 October. Available at: https://deepmind.google/blog/bringing-ai-to-the-next-generation-of-fusion-energy/

128. ATLAS Collaboration (2025) 'Transforming jet flavour tagging at ATLAS', arXiv preprint 2505.19689. Available at: https://arxiv.org/abs/2505.19689

129. Senior, A.W. et al. (2025) 'A scalable and real-time neural decoder for topological quantum codes', arXiv preprint 2512.07737. Available at: https://arxiv.org/abs/2512.07737

130. Yan, G. et al. (2026) 'OmniQEC: discovering practical quantum error-correcting codes by an AI scientist', arXiv preprint 2607.25865, 28 July. Available at: https://arxiv.org/abs/2607.25865

131. Hess, P. (2026) 'Making AI-driven chip design affordable with open agent factories and distributed inference at scale', IBM Research blog, 13 October. Available at: https://research.ibm.com/blog/open-agent-factories-and-distributed-inference-for-chip-design

132. Ivezic, M. (2026) 'IBM's AI agents cut a quantum decoder's estimated latency 52-fold for $26 and 204-fold for $299', *PostQuantum*, 18 October. Available at: https://postquantum.com/engineering-news/ai-chip-design-quantum-decoder/

133. Nellis, S. (2026) 'Microsoft reveals new quantum chip made with AI, says it will have systems by 2029', *Reuters* (via U.S. News & World Report), 2 June. Available at: https://money.usnews.com/investing/news/articles/2026-06-02/microsoft-reveals-new-quantum-chip-made-with-ai-says-it-will-have-systems-by-2029

134. Bolgar, C. (2026) 'Introducing Majorana 2', *Microsoft Source*, 2 June. Available at: https://news.microsoft.com/source/features/innovation/majorana-2-microsoft-discovery-agentic-ai/

135. Quantinuum (2026) 'Quantinuum breaks error-correction world record by 10x', 19 October. Available at: https://www.quantinuum.com/press-releases/quantinuum-breaks-error-correction-world-record-by-10x-validating-its-leading-path-to-fault-tolerant-quantum-computing-by-2029

136. IBM (2026) 'IBM and The University of Chicago demonstrate quantum advantage, establishing trusted quantum computation on logical circuits', 30 July. Available at: https://newsroom.ibm.com/2026-07-30-ibm-and-the-university-of-chicago-demonstrate-quantum-advantage,-establishing-trusted-quantum-computation-on-logical-circuits

137. PostQuantum (2026) 'IBM's three quantum advantage papers', 1 August. Available at: https://postquantum.com/industry-news/ibm-three-quantum-advantage-papers/

138. Neven, H. and Smelyanskiy, V. (2025) 'Our Quantum Echoes algorithm is a big step toward real-world applications for quantum computing', Google blog, 22 October. Available at: https://blog.google/technology/research/quantum-echoes-willow-verifiable-quantum-advantage/

139. Google DeepMind (2023) 'GraphCast: AI model for faster and more accurate global weather forecasting' (*Science*), 14 November. Available at: https://deepmind.google/discover/blog/graphcast-ai-model-for-faster-and-more-accurate-global-weather-forecasting/

140. Google DeepMind (2024) 'GenCast predicts weather and the risks of extreme conditions with state-of-the-art accuracy' (*Nature*), 4 December. Available at: https://deepmind.google/blog/gencast-predicts-weather-and-the-risks-of-extreme-conditions-with-sota-accuracy/

141. ECMWF (2025) 'ECMWF's ensemble AI forecasts become operational', 1 July. Available at: https://www.ecmwf.int/en/about/media-centre/news/2025/ecmwfs-ensemble-ai-forecasts-become-operational

142. Bodnar, C. et al. (2025) 'A foundation model for the Earth system', *Nature*, 641, 1180–1187. doi:10.1038/s41586-025-09005-y. Available at: https://ideas.repec.org/a/nat/nature/v641y2025i8065d10.1038_s41586-025-09005-y.html

143. Kochkov, D. et al. (2024) 'Neural general circulation models for weather and climate', *Nature*, 632, 1060–1066. Available at: https://research.google/pubs/neural-general-circulation-models-for-weather-and-climate/

144. Watt-Meyer, O. et al. (2025) 'ACE2: accurately learning subseasonal to decadal atmospheric variability and forced responses', *npj Climate and Atmospheric Science*, 8. Available at: https://doaj.org/article/95125273a05d4774a836db38118506b7

145. Department of Agriculture, Food and the Marine (2025) 'Minister Heydon and Minister of State Grealish launch €26.2 million VistaMilk Phase II Research Centre', 3 April. Available at: https://www.gov.ie/en/department-of-agriculture-food-and-the-marine/press-releases/minister-heydon-and-minister-of-state-grealish-launch-262-million-vistamilk-phase-ii-research-centre/

146. Argyle, L.P. et al. (2023) 'Out of one, many: using language models to simulate human samples', *Political Analysis*, 31(3), 337–351. Available at: https://www.cambridge.org/core/product/035D7C8A55B237942FB6DBAD7CAA4E49

147. Bisbee, J., Clinton, J.D., Dorff, C., Kenkel, B. and Larson, J.M. (2024) 'Synthetic replacements for human survey data? The perils of large language models', *Political Analysis*, 32(4), 401–416. Available at: https://www.cambridge.org/core/product/B92267DC26195C7F36E63EA04A47D2FE

148. Research Live (2026) 'Synthetic respondents can't replace human measurement, finds Prolific study' (Gordon, A. et al., 'Do richer personas improve LLM survey simulation? A fidelity paradox', preprint), 24 August. Available at: https://www.research-live.com/article/news/synthetic-respondents-cant-replace-human-measurement-finds-prolific-study/id/5152227

149. Binz, M. et al. (2025) 'A foundation model to predict and capture human cognition', *Nature*, 644, 1002–1009. Available at: https://collaborate.princeton.edu/en/publications/a-foundation-model-to-predict-and-capture-human-cognition/

150. University of Bristol (2025) 'Professor Jeffery Bowers calls new AI model of human mind "absurd"', 8 July. Available at: https://www.bristol.ac.uk/psychology-neuroscience/news/2025/121---ai-model-of-human-mind.html

151. Liu, W. and Ding, N. (2026) 'Can Centaur truly simulate human cognition? The fundamental limitation of instruction understanding', *National Science Open*, 5, 20250053. doi:10.1360/nso/20250053. Available at: https://www.nso-journal.org/articles/nso/pdf/2026/01/NSO20250053.pdf

152. Korinek, A. (2023) 'Generative AI for economic research: use cases and implications for economists', *Journal of Economic Literature*, 61(4), 1281–1317. Available at: https://www.aeaweb.org/doi/10.1257/jel.20231736

153. Horton, J.J., Filippas, A. and Manning, B.S. (2023, revised 2026) 'Large language models as simulated economic agents: what can we learn from Homo silicus?', NBER Working Paper 31122. Available at: https://www.nber.org/papers/w31122

154. Hall, P. (2025) 'Google DeepMind's new AI model helps historians interpret Latin inscriptions', *MIT Technology Review*, 23 July. Available at: https://www.technologyreview.com/2025/07/23/1120574/deepmind-ai-aeneas-helps-historians-interpret-latin-inscriptions

155. Flemyng, E. et al. (2025) 'Position statement on artificial intelligence (AI) use in evidence synthesis across Cochrane, the Campbell Collaboration, JBI, and the Collaboration for Environmental Evidence 2025', *JBI Evidence Synthesis*, 23(11), 2162–2166. Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC12594113/

156. Cao, C. et al. (2025) 'Automation of systematic reviews with large language models', medRxiv preprint, 13 June. Available at: https://www.medrxiv.org/content/10.1101/2025.06.13.25329541v1

157. Nobel Prize Outreach (2024) 'The Nobel Prize in Physics 2024 – Press release', 8 October. Available at: https://www.nobelprize.org/prizes/physics/2024/press-release/

158. Stanford Institute for Human-Centered Artificial Intelligence (2026) *The 2026 AI Index Report*, Chapter 1: Research and Development. Available at: https://hai.stanford.edu/assets/files/ai_index_report_2026_chapter_1_research_development.pdf

159. The Decoder (2026) 'Google says 75 percent of its new code is now written by AI', 23 April. Available at: https://the-decoder.com/google-says-75-percent-of-its-new-code-is-now-written-by-ai/

160. Lu, C. et al. (2026) 'Towards end-to-end automation of AI research', *Nature*, 25 March. Available at: https://www.nature.com/articles/s41586-026-10265-5

161. Jenkins, J. (2026) 'The AI Scientist: towards full automation of the research life cycle', Vector Institute, 9 July. Available at: https://vectorinstitute.ai/?p=26786

162. TechCrunch (2026) 'Anthropic warns that it's very close to RSI capability, aka self-building AI', 2 November. Available at: https://techcrunch.com/2026/11/02/anthropic-warns-that-its-very-close-to-rsi-capability-aka-self-building-ai/

163. Takehi, R., Teoh, J., Zhang, A. and Khattab, O. (2026) 'Sparks of weight-space programming: LLMs can natively read and write neural network weights', MIT CSAIL, 2 November. Available at: https://rikiyat.github.io/blog/weight-space-programming/

164. UK AI Security Institute (2026) 'The Alignment Project'. Available at: https://alignmentproject.aisi.gov.uk/

165. OpenAI (2026) 'OpenAI Residency'. Available at: https://openai.com/residency

166. Google Research (2026) 'Visiting Researcher Program'. Available at: https://research.google/outreach/visiting-researcher-program

167. Wiley (2025) 'AI adoption jumps to 84% among researchers as expectations undergo significant reality check', 7 October. Available at: https://newsroom.wiley.com/press-releases/press-release-details/2025/AI-Adoption-Jumps-to-84-Among-Researchers-as-Expectations-Undergo-Significant-Reality-Check/default.aspx

168. Elsevier (2025) 'Elsevier's global survey of 3,000 researchers reveals less than half have time to do research but see AI as transformative', 4 November. Available at: https://www.elsevier.com/about/press-releases/elseviers-global-survey-of-3-000-researchers-reveals-less-than-half-have

169. Springer Nature (2026) 'Early career researchers lead shift towards AI-first discovery', 20 July. Available at: https://group.springernature.com/de/group/media/press-releases/early-career-researchers-lead-shift-towards-ai-first-discovery/52988132

170. Office for Students (2026) *Financial sustainability of higher education providers in England 2026*, OfS 2026.21, 14 May (corrected edition, 15 June). Available at: https://www.officeforstudents.org.uk/media/p0dd3sdq/ofs-2026_21-updated-15-6-26.pdf

171. Hermle, J., Jones, B., Althoff, L. et al. (2026) 'AI for scientific discovery: how researchers use Claude', Anthropic, 5 November. Available at: https://www.anthropic.com/research/how-researchers-use-claude

172. OpenAI (2026) 'Science at the frontier', 16 November. Available at: https://openai.com/index/science-at-the-frontier/

173. National Institutes of Health (2025) 'Supplemental guidance to the 2024 NIH Grants Policy Statement: indirect cost rates', NOT-OD-25-068, 7 February. Available at: https://grants.nih.gov/grants/guide/notice-files/not-od-25-068.html

174. Schneider, S.L. et al. (2020) *2018 FDP Faculty Workload Survey: Report of Primary Findings*. Washington, DC: Federal Demonstration Partnership. Available at: https://thefdp.org/wp-content/uploads/FDP-FWS-2018-Primary-Report.pdf

175. Hao, Q., Xu, F., Li, Y. and Evans, J. (2026) 'Artificial intelligence tools expand scientists' impact but contract science's focus', *Nature*, doi:10.1038/s41586-025-09922-y. Available at: https://doi.org/10.1038/s41586-025-09922-y; preprint: https://arxiv.org/abs/2412.07727

176. Kapoor, S. and Narayanan, A. (2023) 'Leakage and the reproducibility crisis in machine-learning-based science', *Patterns*, 4(9), 100804. Available at: https://collaborate.princeton.edu/en/publications/leakage-and-the-reproducibility-crisis-in-machine-learning-based-/

177. Kapoor, S. et al. (2024) 'REFORMS: consensus-based recommendations for machine-learning-based science', *Science Advances*. Available at: https://reforms.cs.princeton.edu/

178. MIT Department of Economics (2025) 'Assuring an accurate research record', 16 May. Available at: https://economics.mit.edu/news/assuring-accurate-research-record

179. Davis, M.F. and Torrence, R. (2026) 'Lila Sciences said in talks for funds at $8.5 billion valuation', *Bloomberg News* (via Bloomberg Law), 3 June. Available at: https://news.bloomberglaw.com/private-equity/lila-sciences-said-in-talks-for-funds-at-8-5-billion-valuation

180. Martin, I. (2026) 'Former OpenAI researcher to raise $500 million for AI science startup', *Forbes* (Forbes Australia), 7 May. Available at: https://www.forbes.com.au/news/innovation/former-openai-researcher-to-raise-500-million-for-ai-science-startup/

181. Isomorphic Labs (2026) 'Isomorphic Labs secures $2.1 billion funding to scale its AI drug design engine', 12 May. Available at: https://www.isomorphiclabs.com/press/isomorphic-labs-funding

182. Edison Scientific (2026) *Edison Platform documentation: Overview*. Available at: https://docs.edisonscientific.com/

183. Lin, F. (2026) 'Boltz PBC launches with $28M to democratize AI platforms for drug discovery', *GEN – Genetic Engineering & Biotechnology News*, 9 January. Available at: https://www.genengnews.com/topics/artificial-intelligence/boltz-pbc-launches-with-28m-to-democratize-ai-platforms-for-drug-discovery

184. Xaira Therapeutics (2024) 'Xaira Therapeutics launches to deliver transformative medicines by advancing and harnessing AI for drug discovery and development', press release, 23 April. Available at: https://www.businesswire.com/news/home/20240423707240/en/Xaira-Therapeutics-Launches-to-Deliver-Transformative-Medicines-by-Advancing-and-Harnessing-AI-for-Drug-Discovery-and-Development

185. Tech.eu (2026) 'Jeff Bezos and Sovereign AI back CuspAI in $450M raise', 20 July. Available at: https://tech.eu/2026/07/20/jeff-bezos-and-sovereign-ai-back-cuspai-in-450m-raise

186. US Department of Energy (2026) 'U.S. Department of Energy launches the Genesis Open Models Initiative', 7 August. Available at: https://www.energy.gov/node/4859769

187. Orrick (2026) 'Orrick advises Mistral in its €3B Series D funding round at €21B post-money valuation', 8 September. Available at: https://www.orrick.com/en/news/2026/09/orrick-advises-mistral-in-its-3b-series-d-at-21b-post-money-valuation

188. State Council of the People's Republic of China (2025) 'Opinions of the State Council on deepening the implementation of the "Artificial Intelligence+" initiative', 21 August (released 26 August), translated by CSET. Available at: https://cset.georgetown.edu/wp-content/uploads/t0652_AI_plus_opinions_EN.pdf

189. The White House (2026) 'Trump Administration announces more than $5 billion for the Genesis Mission, a national mission on AI for science', 22 July. Available at: https://www.whitehouse.gov/releases/2026/07/45502/

190. US Department of Energy Office of Science (2026) 'Energy Department announces new Genesis Mission awards to advance super intelligence for science' [GovDelivery bulletin], 9 October. Available at: https://content.govdelivery.com/accounts/USDOEOS/bulletins/42e95ea

191. US Department of Energy (2026) 'Secretary of Energy Chris Wright announces first Genesis Mission projects selected to accelerate...', 22 July. Available at: https://www.energy.gov/articles/secretary-energy-chris-wright-announces-first-genesis-mission-projects-selected-accelerate

192. San Diego Supercomputer Center (2026) 'SDSC to lead NSF NAIRR Operations Center', 1 September. Available at: https://www.sdsc.edu/news/2026/PR20260901-NAIRR-OC.html

193. Kelley, A. (2026) 'NSF launches a new AI research operations center', *Nextgov/FCW*, 2 September. Available at: https://www.nextgov.com/artificial-intelligence/2026/09/nsf-launches-new-ai-research-operations-center/415786/

194. National Science Foundation (2026) [Programmable Cloud Laboratories Test Bed awards] (GovDelivery bulletin), 23 July. Available at: https://content.govdelivery.com/accounts/USNSF/bulletins/421a483

195. National Science Foundation (2025) 'NSF to invest in new national network of AI programmable cloud laboratories', 5 August. Available at: https://www.nsf.gov/news/nsf-invest-new-national-network-ai-programmable-cloud

196. Department for Science, Innovation and Technology (2025) *UK Compute Roadmap*, 17 July (updated 23 April 2026). Available at: https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

197. European Commission (2025) *European Strategy for AI in Science*, 8 October. Available at: https://research-and-innovation.ec.europa.eu/strategy/strategy-research-and-innovation/our-digital-future/european-ai-science-strategy_en

198. EuroHPC Joint Undertaking (2026) *Invitation to tender to procure CASPIr supercomputer*, 27 March. Available at: https://www.eurohpc-ju.europa.eu/invitation-tender-procure-caspir-supercomputer-2026-03-27_en

199. EuroHPC Joint Undertaking (2026) 'EuroHPC Joint Undertaking launches AI Gigafactories call', 30 July. Available at: https://eurohpc-ju.europa.eu/eurohpc-joint-undertaking-launches-ai-gigafactories-call-2026-07-30_en

200. European Commission (2026) 'Commission launches RAISE pledging initiative to boost AI support for researchers across Europe', press release IP/26/2175, 15 October. Available at: https://ec.europa.eu/commission/presscorner/detail/en/ip_26_2175

201. Spenner, E.A. (2026) 'Harvard weighs $25 million investment to double university-wide AI computing capacity', *The Harvard Crimson*, 5 October. Available at: https://thecrimson.com/article/2026/10/5/harvard-ai-computing-investment

202. Princeton University (2024) 'Princeton invests in new 300-GPU cluster for academic AI research', 15 March. Available at: https://ai.princeton.edu/news/2024/princeton-invests-new-300-gpu-cluster-academic-ai-research

203. Mid Sweden University (2026) 'New supercomputer will strengthen research and education at Mid Sweden University', 27 August. Available at: https://www.miun.se/en/Research/research-centers/stc/news-from-stc-research-center-eng/News-archive/2026-8/New-supercomputer-will-strengthen-research-and-education-at-Mid-Sweden-University/

204. UKRI (2026) 'Isambard-AI and Zenith AIRR supercomputers: rapid access route', updated 15 October. Available at: https://www.ukri.org/opportunity/isambard-ai-and-zenith-airr-supercomputers-rapid-access-route/

205. CommonAI CIC (2026) *Home*. Available at: https://commonai.org/

206. PR Newswire (2017) 'Sci-Tech Daresbury and IBM Research Collaborative Campus at Hartree Centre Receives Royal Seal of Approval', 31 May. Available at: https://prnewswire.co.uk/news-releases/sci-tech-daresbury-and-ibm-research-collaborative-campus-at-hartree-centre-receives-royal-seal-of-625530244.html

207. UKSPA (2018) 'Collaboration drives £17.3m boost for Sci-Tech Daresbury companies', 11 July. Available at: https://www.ukspa.org.uk/collaboration-drives-17-3m-boost-for-sci-tech-daresbury-companies/

208. Research Professional News (2019) 'Daresbury welcomes 150th tenant to Cheshire science campus', December. Available at: https://www.researchprofessionalnews.com/rr-news-uk-innovation-2019-12-daresbury-welcomes-150th-tenant-to-cheshire-science-campus/

209. Innovate UK Business Connect (2025) 'Liverpool Innovation Investment Summit highlights regional scale-up potential', 10 June. Available at: https://iuk-business-connect.org.uk/news/liverpool-innovation-investment-summit-highlights-regional-scale-up-potential/

210. Cuff, J. (2018) 'The practice of collaboration between industry and HPC centers', *The Next Platform*, 10 July. Available at: https://www.nextplatform.com/2018/07/10/the-practice-of-collaboration-between-industry-and-hpc-centers

211. CSC – IT Center for Science (2025) 'Kajaani data center ecosystem strengthens its position – CSC an active partner in the collaboration', 19 March. Available at: https://csc.fi/en/news/kajaani-data-center-ecosystem-strengthens-its-position/

212. LUMI (2025) 'LUMI AI Factory first services', 6 October. Available at: https://www.lumi-supercomputer.eu/lumi-ai-factory-first-services/

213. EuroHPC Joint Undertaking (2026) 'EuroHPC JU signs contract to deploy LUMI-AI supercomputer', 31 August. Available at: https://www.eurohpc-ju.europa.eu/eurohpc-ju-signs-contract-deploy-lumi-ai-supercomputer-2026-08-31_en

214. City of Kajaani (2025) 'XTX Markets plans to invest over €1bn in large-scale data centre project in Kajaani, Finland', 22 January. Available at: https://kajaani.fi/en/news/xtx-markets-plans-to-invest-over-e1bn-in-large-scale-data-centre-project-in-kajaani-finland/

215. STFC Hartree Centre (2025) *Hartree Highlights 2024–25*, September. Available at: https://www.hartree.stfc.ac.uk/wp-content/uploads/2025/09/Hartree_Highlights_2024-25.pdf

216. Unilever UK (2026) 'How Unilever and the University of Liverpool are turning world-class science into innovations that improve lives', 11 June. Available at: https://www.unilever.co.uk/news/2026/how-unilever-and-the-university-of-liverpool-are-turning-worldclass-science-into-innovations-that-improve-lives/

217. University of Toronto (2023) 'Acceleration Consortium' [CFREF award announcement], Defy Gravity campaign news, 28 April. Available at: https://defygravitycampaign.utoronto.ca/news-and-stories/acceleration-consortium-cfref/

218. Acceleration Consortium (2025) 'A year in review', 23 December. Available at: https://accelerationconsortium.substack.com/p/a-year-in-review-96d

219. Anthropic (2026) 'Measurements for understanding the pace of AI development', figures as of August 2026. Available at: https://www.anthropic.com/institute/measuring-pace-of-ai-development

220. OpenAI (2026) 'Research acceleration: the jagged frontier', 18 November. Available at: https://openai.com/index/research-acceleration-jagged-frontier/

221. METR (2025) 'Measuring the impact of early-2025 AI on experienced open-source developer productivity', 10 July. Available at: https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/

222. METR (2026) 'We are changing our developer productivity experiment design', 24 February. Available at: https://metr.org/blog/2026-02-24-uplift-update/

223. Demirer, M., Musolff, L. and Yang, L. (2026) 'Writing code vs. shipping code: productivity effects across generations of AI coding tools', NBER Working Paper 35275, May, revised September. Available at: https://www.nber.org/papers/w35275

224. Anthropic (2026) 'How AI assistance impacts the formation of coding skills', 29 January. Available at: https://www.anthropic.com/research/AI-assistance-coding-skills

225. Cui, Z.K., Demirer, M., Jaffe, S., Musolff, L., Peng, S. and Salz, T. (2025) 'The effects of generative AI on high-skilled work: evidence from three field experiments with software developers', *Management Science*. Available at: https://doi.org/10.1287/mnsc.2025.00535

226. Harmon, J. (2026) 'Argonne's upgraded X-ray source being linked with nation's top supercomputers to handle increased data flow', Argonne National Laboratory, 26 January. Available at: https://www.anl.gov/article/argonnes-upgraded-xray-source-being-linked-with-nations-top-supercomputers-to-handle-increased-data

227. Diamond Light Source (2026) 'OpenBind's first data and model release marks a milestone for AI enabled drug discovery', 6 May. Available at: https://www.diamond.ac.uk/default/Home/News/LatestNews/2026/OpenBind-s-first-data-and-model-release-marks-a-milestone-for-AI-enabled-drug-discovery.html

228. Vithayapalert, N. and Grisoni, F. (2026) 'Towards generalizable protein-ligand co-folding with ACER', *bioRxiv*, June; presented at ICML 2026. doi:10.64898/2026.06.02.728568. Available at: https://doi.org/10.64898/2026.06.02.728568

229. Mao, X. et al. (2026) 'Benchmarking virtual cell models for in-the-wild perturbation response', *arXiv*, 2604.27646, 30 April. Available at: https://arxiv.org/abs/2604.27646

230. Jishnu, A. and Ansari, E. (2026) 'AI memory boom squeezes legacy DRAM supply, pushing prices higher', S&P Global Market Intelligence, 28 January. Available at: https://www.spglobal.com/market-intelligence/en/news-insights/research/2026/01/ai-memory-boom-squeezes-legacy-dram-supply-pushing-prices-higher

231. Anthropic (2026) 'Statement on the US government directive to suspend access to Fable 5 and Mythos 5', 12 June. Available at: https://www.anthropic.com/news/fable-mythos-access

232. Google DeepMind (2025) 'Genie 3: a new frontier for world models', 5 August. Available at: https://deepmind.google/discover/blog/genie-3-a-new-frontier-for-world-models/

233. NVIDIA (2025) 'NVIDIA launches Cosmos world foundation model platform to accelerate physical AI development', 6 January. Available at: https://nvidianews.nvidia.com/news/nvidia-launches-cosmos-world-foundation-model-platform-to-accelerate-physical-ai-development

234. Deutscher, M. (2026) 'Yann LeCun's new startup AMI Labs raises $1.03B to train world models', *SiliconANGLE*, 10 March. Available at: https://siliconangle.com/2026/03/10/yann-lecuns-new-startup-ami-labs-raises-1-03b-train-world-models/

235. Cunningham et al. (2026) 'The economics of recursive self-improvement', arXiv preprint 2609.15802. Available at: https://arxiv.org/abs/2609.15802

236. Kučinskas, S. and Rosenberg, J. (2026) 'How close is recursive self-improvement? A review of the evidence', Forecasting Research Institute, 10 November. Available at: https://forecastingresearch.org/research/recursive-self-improvement-v1

237. Ropes & Gray (2025) 'Congress passes final tax package: key provisions for tax-exempt organizations', July. Available at: https://www.ropesgray.com/en/insights/alerts/2025/07/congress-passes-final-tax-package-key-provisions-for-tax-exempt-organizations

238. Google (2025) 'Gemini app now available with enterprise-grade data protections for all Google Workspace for Education users', 28 March. Available at: https://blog.google/outreach-initiatives/education/gemini-enterprise-grade-data-protection

239. University of Cambridge Information Services (2024) 'Copilot', IT Help and Support, 28 March. Available at: https://help.uis.cam.ac.uk/node/15002

240. STM Association (2026) 'Uptake of open access', OA Dashboard. Available at: https://stm-assoc.org/oa-dashboard/oa-dashboard/uptake-of-open-access/

241. Directive (EU) 2019/790 on copyright and related rights in the Digital Single Market, Article 3, 'Text and data mining for the purposes of scientific research'. Available at: https://eur-lex.europa.eu/eli/dir/2019/790/oj

242. Copyright, Designs and Patents Act 1988, section 29A, 'Copies for text and data analysis for non-commercial research'. Available at: https://www.legislation.gov.uk/ukpga/1988/48/section/29A

243. Digital Health and Care Northern Ireland (n.d.) 'encompass'. Available at: https://dhcni.hscni.net/digital-portfolio/encompass/

244. Health Data Research UK (n.d.) 'Northern Ireland Honest Broker Service (Northern Ireland Health and Social Care)'. Available at: https://www.hdruk.ac.uk/organisations/northern-ireland-honest-broker-service-northern-ireland-health-and-social-care/

245. Enterprise Ireland Technology Gateway Network (2025) 'What are Gateways?'. Available at: https://technologygateway.ie/network/

246. UK Research and Innovation (2025) 'Knowledge Transfer Partnership guidance: who can take part', Innovate UK, updated 27 October. Available at: https://www.ukri.org/councils/innovate-uk/guidance-for-applicants/guidance-for-specific-funds/knowledge-transfer-partnership-guidance/who-can-take-part/

247. National Science Foundation (2026) 'About IUCRC', Industry–University Cooperative Research Centers Program. Available at: https://iucrc.nsf.gov/about

248. High Value Manufacturing Catapult (2026) 'What is a Catapult?'. Available at: https://hvm.catapult.org.uk/who-we-are/what-is-a-catapult/

249. Ridgway, K. (2026) 'Universities built the Catapults. Now they should release them', *Times Higher Education*, 15 October. Available at: https://www.timeshighereducation.com/node/745511

250. National Institute of Standards and Technology (2026) 'NIST Manufacturing Extension Partnership (MEP)', updated 22 June. Available at: https://www.nist.gov/mep

251. Enterprise Ireland (2026) 'European Digital Innovation Hubs'. Available at: https://www.enterprise-ireland.com/en/supports/become-more-innovative/european-digital-innovation-hubs

252. University College Dublin (2026) 'CeADAR secures €5.7 million for phase two of European Digital Innovation Hub programme', 9 January. Available at: https://www.ucd.ie/innovation/news-and-events/2026/ceadar-phase-2-edih-funding

253. Oxford University Innovation (2026) 'University of Oxford confirmed as UK and European leader for spinouts in key report', 29 June. Available at: https://innovation.ox.ac.uk/news/university-oxford-confirmed-uk-and-european-leader-spinouts-key-report

254. University of Oxford (2022) 'Oxford Science Enterprises raises £250 million to support University innovations', 18 July. Available at: https://www.ox.ac.uk/news/2022-07-18-oxford-science-enterprises-raises-250-million-support-university-innovations

255. ai@cam (2026) *ai@cam Annual Report 2025–2026*, University of Cambridge. Available at: https://www.ai.cam.ac.uk/reports/ai-cam-annual-report-2025-2026/

256. National Technology News (2025) 'CIC to invest "at least" £100m in Cambridge University spinouts'. Available at: https://nationaltechnology.co.uk/CIC_To_Invest_At_Least_100m_In_Cambridge_University_Spinouts.php

257. UK Supreme Court (2023) *Thaler v Comptroller-General of Patents, Designs and Trade Marks* [2023] UKSC 49, 20 December (press summary). Available at: https://caselaw.nationalarchives.gov.uk/uksc/2023/49/press-summary

258. European Patent Office (2021) 'Legal Board of Appeal decision J 8/20 (DABUS)', 21 December. Available at: https://www.epo.org/en/law-and-practice/boards-of-appeal/communications/press-communique-decisions-j-820-and-j-920-legal

259. US Court of Appeals for the Federal Circuit (2022) *Thaler v. Vidal*, No. 2021-2347, 43 F.4th 1207, 5 August. Available at: https://www.cafc.uscourts.gov/opinions-orders/21-2347.OPINION.8-5-2022_1988142.pdf

260. US Patent and Trademark Office (2025) 'Revised inventorship guidance for AI-assisted inventions', *Federal Register*, 90 FR 54636, 28 November. Available at: https://www.federalregister.gov/documents/2025/11/28/2025-21457/revised-inventorship-guidance-for-ai-assisted-inventions

261. US Copyright Office (2025) *Copyright and Artificial Intelligence, Part 2: Copyrightability*, 29 January. Available at: https://www.copyright.gov/ai/

262. UK Government (2026) *Report on Copyright and Artificial Intelligence*, March. Available at: https://assets.publishing.service.gov.uk/media/69ba692226909a14239612e4/CP2602959_-_Report_on_Copyright_and_Artificial_Intelligence_web.pdf

263. Hugging Face (2022) 'OpenRAIL: towards open and responsible AI licensing frameworks', 31 August. Available at: https://huggingface.co/blog/open_rail

264. Google DeepMind (2024) *AlphaFold 3 model parameters terms of use*. Available at: https://github.com/google-deepmind/alphafold3/blob/main/WEIGHTS_TERMS_OF_USE.md

265. European Commission (2025) 'The General-Purpose AI Code of Practice', published 10 July. Available at: https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai

266. Keystone Law (2026) 'Emotional Perception: UK Supreme Court sets out a new approach to patentability for computer implemented inventions'. Available at: https://keystonelaw.com/keynotes/emotional-perception-uk-supreme-court-sets-out-a-new-approach-to-patentability-for-computer-implemented-inventions/

267. STFC Hartree Centre (2023) *HNCDI EXCELERATE IP Template Agreement*, June. Available at: https://www.hartree.stfc.ac.uk/wp-content/uploads/2023/07/HNCDI-EXCELERATE-IP-Template-Agreement-June-2023.pdf

268. Structural Genomics Consortium (n.d.) 'About'. Available at: https://www.thesgc.org/about

269. Knowledge Transfer Ireland (2019) *National IP Protocol*. Available at: https://www.enterprise-ireland.com/documents/ireland-s-national-ip-protocol-2019-en-253454.pdf

270. Intellectual Property Office (2022) 'Lambert Toolkit', updated 5 January. Available at: https://www.gov.uk/lambert-toolkit

271. Uttley, L. (2026) *Review and Analysis of UK Institutional Guidance for Artificial Intelligence Use in Research: Executive Summary*. UK Committee on Research Integrity, March. Available at: https://ukcori.org/wp-content/uploads/2026/03/Executive-Summary-Review-and-Analysis-of-UK-Institutional-Guidance-for-Artificial-Intelligence-Use-in-Research.pdf

272. Regulation (EU) 2016/679 (General Data Protection Regulation), Article 35, 'Data protection impact assessment'. Available at: https://eur-lex.europa.eu/eli/reg/2016/679/oj

273. US Department of Health and Human Services, 'Protection of Human Subjects', 45 CFR Part 46 (the Common Rule). Available at: https://www.ecfr.gov/current/title-45/subtitle-A/subchapter-A/part-46

274. US Department of Health and Human Services, 'Research' (HIPAA Privacy Rule guidance). Available at: https://www.hhs.gov/hipaa/for-professionals/special-topics/research/index.html

275. US Department of Education, 'Family Educational Rights and Privacy Act (FERPA)', 20 U.S.C. § 1232g; 34 CFR Part 99. Available at: https://studentprivacy.ed.gov/ferpa

276. Regulation (EU) 2024/1689 (Artificial Intelligence Act), Articles 2(6), 2(8) and 4. Available at: https://eur-lex.europa.eu/eli/reg/2024/1689/oj

277. Regulation (EU) 2026/1744 amending Regulation (EU) 2024/1689 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), *Official Journal of the European Union*, L series, 24 July 2026. Available at: https://eur-lex.europa.eu/eli/reg/2026/1744/oj

278. Lewis Silkin (2026) 'The Digital Omnibus on AI enters into force today', 27 July. Available at: https://www.lewissilkin.com/insights/2026/07/27/the-digital-omnibus-on-ai-enters-into-force-today

279. Cloud Security Alliance (2026) 'State AI laws take hold as federal preemption stalls', CSA research note, April. Available at: https://labs.cloudsecurityalliance.org/research/csa-research-note-us-ai-regulation-preemption-compliance-202/

280. Commission Recommendation (EU) 2021/1700 of 15 September 2021 on internal compliance programmes for controls of research involving dual-use items under Regulation (EU) 2021/821. Available at: https://eur-lex.europa.eu/eli/reco/2021/1700/oj

281. National Protective Security Authority (2026) 'Trusted Research'. Available at: https://www.npsa.gov.uk/trusted-research

282. Infosecurity Magazine (2026) 'UK: education sector faces surge in cyber breaches despite stable national threat levels', 30 April, reporting the Cyber Security Breaches Survey 2025/2026. Available at: https://infosecurity-magazine.com/news/uk-education-sector-faces-surge-in

283. Higher Education Authority and Research Ireland (2026) *INSPIRE Research Infrastructure Programme Phase One: Call Document*, May. Available at: https://hea.ie/assets/uploads/2026/02/INSPIRE-Phase-One-Call-Document.pdf

284. Council of the European Union (2021) 'Council Regulation (EU) 2021/1173 establishing the European High Performance Computing Joint Undertaking', 13 July, Articles 11(2) and 14(2) (since amended). Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32021R1173

285. National Science Foundation (2023) *Major Research Instrumentation (MRI) Program* (solicitation NSF 23-519). Available at: https://www.nsf.gov/funding/opportunities/mri-major-research-instrumentation-program/nsf23-519/solicitation

286. University at Buffalo (2025) 'Empire AI supercomputing gets funding boost', *UBNow*, 30 June. Available at: https://www.buffalo.edu/ubnow/stories/2025/06/empire-ai-supercomputing-funding.html

287. IntuitionLabs (2026) 'Data center GPU pricing 2026: the full AI pricing index', updated 11 October. Available at: https://intuitionlabs.ai/articles/data-center-gpu-pricing-2026

288. NVIDIA (n.d.) 'Introduction to the NVIDIA DGX H100 System', *NVIDIA DGX H100/H200 User Guide*. Available at: https://docs.nvidia.com/dgx/dgxh100-user-guide/introduction-to-dgxh100.html

289. University of Toronto (2025) 'With AI and robotics, U of T students build "self-driving" lab for less than $500', 3 July. Available at: https://www.utoronto.ca/news/ai-and-robotics-u-t-students-build-self-driving-lab-less-500

290. Gaulin, M. (n.d.) 'How self-driving labs can cut automation costs by 90 percent', *Lab Manager*. Available at: https://www.labmanager.com/how-self-driving-labs-can-cut-automation-costs-by-90-percent-35276

291. Brazil, R. (2026) '"Self-driving" labs power up', *Nature*, 652, 262–264, 2 April. Available at: https://media.nature.com/original/magazine-assets/d41586-026-00974-2/d41586-026-00974-2.pdf

292. Research Ireland (2025) *Research Ireland Infrastructure Programme 2025: Call for Submission of Proposals*, 26 May. Available at: https://media.researchireland.ie/wp-content/uploads/2026/02/Research-Ireland_Infrastructure-Programme-2025_Call-Document.pdf

293. EuroHPC Joint Undertaking (2025) 'The way is open to building a EuroHPC world-class supercomputer in Ireland', 13 October. Available at: https://www.eurohpc-ju.europa.eu/way-open-building-eurohpc-world-class-supercomputer-ireland-2025-10-13_en

294. National Science Foundation (2024) *Advanced Computing Systems & Services: Adapting to the Rapid Evolution of Science and Engineering Research* (solicitation NSF 24-583). Available at: https://www.nsf.gov/funding/opportunities/advanced-computing-systems-services-adapting-rapid-evolution/nsf24-583/solicitation

295. UKRI (2025) 'Phase 2 National Compute Resource: expression of interest', 21 October. Available at: https://engagementhub.ukri.org/ukri-infrastructure/phase2-national-compute-resource-expression-of-int/

296. European Commission (2022) 'Communication from the Commission – Framework for State aid for research and development and innovation' (2022/C 414/01), *Official Journal of the European Union*, C 414, 28 October, paragraph 21. Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52022XC1028(03)

297. Research Ireland (2026) *Research Ireland Infrastructure Programme 2025: FAQs*, question 22. Available at: https://media.researchireland.ie/wp-content/uploads/2026/02/Research-Ireland-Infrastructure-Programme-2025-FAQs.pdf

298. Department for Business and Trade (2025) *UK subsidy control regime: statutory guidance*, first published 11 November 2022, updated 4 August 2025. Available at: https://www.gov.uk/government/publications/uk-subsidy-control-statutory-guidance

299. Farrer & Co (2024) 'Subsidy control: what does it mean for universities and research institutions?', February. Available at: https://www.farrer.co.uk/news-and-insights/subsidy-control-what-does-it-mean-for-universities-and-research-institutions/

300. Department for Business and Trade (n.d.) 'Guidance on the scope and application of Article 10 of the Windsor Framework', *Complying with the UK's international obligations on subsidy control*. Available at: https://www.gov.uk/government/publications/complying-with-the-uks-international-obligations-on-subsidy-control-guidance-for-public-authorities/guidance-on-the-scope-and-application-of-article-10-of-the-windsor-framework

301. US Code of Federal Regulations, 2 CFR § 200.313, 'Equipment'. Available at: https://www.law.cornell.edu/cfr/text/2/200.313

302. UIDP (2021) 'Tax exempt bonds and research contracts', session materials reproducing COGR's *Tax Issues in Industry-University Relationships* and IRS Rev. Proc. 2007-47, 15 April. Available at: https://uidp.org/wp-content/uploads/2021/04/Tax-Exempt-Bonds-and-Research-Contracts.pdf

303. European Commission (2026) 'Pilot access schemes to Technology Infrastructures for European startups, scaleups and innovative SMEs (CSA)', Horizon Europe topic HORIZON-CL4-2027-01-MAT-PROD-47, deadline 2 February 2027, EU Funding & Tenders Portal. Available at: https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/HORIZON-CL4-2027-01-MAT-PROD-47

304. Koundinya, S. (2026) 'Why NVIDIA's six-year-old GPU is still making money', *Analytics India Magazine*, 17 August. Available at: https://analyticsindiamag.com/ai-features/why-nvidias-six-year-old-gpu-is-still-making-money

305. Fields, K. (2026) 'UK supercomputer gets £36m boost and AMD's latest AI chips', *The Stack*, 26 January. Available at: https://www.thestack.technology/uk-supercomputer-gets-36m-boost-and-amds-latest-ai-chips/

306. DatacenterDynamics (2026) 'Cambridge University launches £36m Zenith supercomputer', June. Available at: https://www.datacenterdynamics.com/en/news/cambridge-university-launches-36m-zenith-supercomputer/

307. Sharwood, S. (2026) 'Google to sell its TPUs to some customers, who also fancy big-G GPUs', *The Register*, 30 April. Available at: https://www.theregister.com/off-prem/2026/04/30/google-to-sell-its-tpus-to-some-customers/5220078

308. Google for Developers (2026) 'TorchTPU: running PyTorch natively on TPUs at Google scale', 7 April. Available at: https://developers.googleblog.com/torchtpu-running-pytorch-natively-on-tpus-at-google-scale/

309. Google Research (2026) 'TPU Research Cloud'. Available at: https://sites.research.google/trc/about/

310. University College Dublin (2025) '€750 million investment package to transform Ireland's research infrastructure', 27 November. Available at: https://www.ucd.ie/newsandopinion/news/2025/november/27/750millioninvestmentpackagetotransformirelandsresearchinfrastructure/

311. ICHEC (2025) 'The EU and Ireland's Department of FHERIS announce the AI Factory Antenna in Ireland (AIF IRL-Antenna)', 13 October. Available at: https://www.ichec.ie/node/1114

312. Research Ireland (2026) *Government announces €460 million investment for seven new Rinn Research Ireland centres*, 10 June. Available at: https://www.researchireland.ie/news/rinn-network/

313. EOSC Association (2024) 'Building the "people ecosystem" in Ireland: interview with Roberto Sabatino of HEAnet', 4 March. Available at: https://eosc.eu/news/building-the-people-ecosystem-in-ireland-interview-with-roberto-sabatino-of-heanet

314. Sabatino, R. et al. (2025) *Implementing the European Open Science Cloud (EOSC) in Ireland*, version 2, Zenodo, 21 July. Available at: https://zenodo.org/records/16266119

315. Research Ireland (n.d.) 'US-Ireland R&D Partnership Programme'. Available at: https://www.researchireland.ie/funding/us-ireland/

316. LUMI AI Factory (2026) *EuroHPC JU access calls for AI Factories*, 5 March. Available at: https://lumi-supercomputer.eu/access-calls-for-ai-factories/

317. EuroHPC Joint Undertaking (2026) 'Access policy and FAQ' (accessed 25 October 2026). Available at: https://eurohpc-ju.europa.eu/access-our-supercomputers/access-policy-and-faq_en

318. EURAXESS (2025) '€100 million AI science pilot projects under Horizon Europe 2026-27', 8 October. Available at: https://euraxess.ec.europa.eu/worldwide/india/news/eu100-million-ai-science-pilot-projects-under-horizon-europe-2026-27

319. EuroHPC Joint Undertaking (2026) 'Call for tenders for the selection of Artificial Intelligence Gigafactory consortia and the establishment of AI Gigafactories', updated 5 November. Available at: https://www.eurohpc-ju.europa.eu/call-tenders-selection-artificial-intelligence-gigafactory-consortia-and-establishment-ai_en

320. Gaudron, J.-M. (2025) 'New Horizon Europe plan nearly doubles funding', Luxinnovation, 29 July. Available at: https://luxinnovation.lu/news/new-horizon-europe-plan-nearly-doubles-funding

321. Zubașcu, F. (2026) 'Partial agreement on the next Horizon Europe sets Council on collision course with Parliament in upcoming negotiations', *Science|Business*, 25 June. Available at: https://sciencebusiness.net/news/planning-fp10/council-pushes-greater-say-shaping-fp10-priorities

322. EuroHPC Joint Undertaking (2024) 'United Kingdom joins EuroHPC Joint Undertaking', 12 May. Available at: https://eurohpc-ju.europa.eu/united-kingdom-joins-eurohpc-joint-undertaking-2024-05-12_en

323. EPCC (n.d.) 'UK AI Factory Antenna: Pilot-UKAIFA'. Available at: https://www.epcc.ed.ac.uk/uk-ai-factory-antenna-pilot-ukaifa

324. Milmo, D. (2026) 'UK can't rely on AI that could be switched off, says Alan Turing Institute chief', *The Guardian*, 10 October, as reported by The Next Web. Available at: https://thenextweb.com/news/alan-turing-institute-williamson-sovereign-ai-uk

325. UKRI (2026) 'UKRI AI strategy makes bold choices where UK can lead the world', 19 February. Available at: https://www.ukri.org/news/ukri-ai-strategy-makes-bold-choices-where-uk-can-lead-the-world

326. Department for Science, Innovation and Technology (2025) 'Sovereign AI open call: autonomous labs – preliminary market engagement', 24 November. Available at: https://www.gov.uk/government/publications/sovereign-ai-open-call-autonomous-labs/sovereign-ai-open-call-autonomous-labs-preliminary-market-engagement

327. UK Research and Innovation (2026) 'Administrative Data Research UK (ADR UK)', area of investment and support. Available at: https://www.ukri.org/what-we-do/browse-our-areas-of-investment-and-support/administrative-data-research-uk-adr-uk/

328. University of Bristol (2026) 'Isambard-AI supercomputer powers £500m UK Sovereign AI fund', 20 April. Available at: https://www.bristol.ac.uk/research/centres/bristol-supercomputing/articles/2026/isambard-ai-supercomputer-powers-500m-uk-sovereign-ai-fund.html

329. US Department of Energy (2026) 'Energy Department launches Genesis Mission Graduate Fellowship to advance American scientific leadership', 8 October. Available at: https://www.energy.gov/articles/energy-department-launches-genesis-mission-graduate-fellowship-advance-american-scientific

330. Graham, E. (2026) 'White House unveils billions in investments for Genesis Mission', *Nextgov/FCW*, 8 October. Available at: https://www.nextgov.com/artificial-intelligence/2026/10/white-house-unveils-billions-investments-genesis-mission/416517/

331. US Department of Energy (2026) 'FACT SHEET: The Energy Department is ushering in a golden age of American science', 8 October. Available at: https://www.energy.gov/node/4860596

332. Krzaczyński, R. (2025) 'OpenEuroLLM: Europe's new initiative for open-source AI development', *InfoQ*, 7 February. Available at: https://www.infoq.com/news/2025/02/open-euro-llm

333. US Department of Energy (2026) 'Department of Energy seeks input on advancing the AI science and engineering workforce', 16 January. Available at: https://www.energy.gov/node/4855799

334. Research Ireland Centre for Research Training in Artificial Intelligence (2026) 'CRT-AI'. Available at: https://crt-ai.ie/

335. The White House (2025) 'Restoring Gold Standard Science' (Executive Order 14303), 23 May. Available at: https://www.whitehouse.gov/presidential-actions/2025/05/restoring-gold-standard-science/

336. ERA Portal Austria (2026) 'Commission publishes update of ERA Living Guidelines on responsible use of generative AI in research', 11 May. Available at: https://community.era.gv.at/news-items/commission-publishes-update-of-era-living-guidelines-on-responsible-use-of-generative-ai-in-research/

337. Congressional Research Service (2026) 'Universities and indirect costs for federally funded research' (R48540), updated 29 January. Available at: https://www.everycrsreport.com/reports/R48540.html

338. US Department of Energy (2025) 'Energy Department announces collaboration agreements with 24 organizations to advance the Genesis Mission', 18 December. Available at: https://www.energy.gov/node/4855568

339. Rock-Torcivia, J. (2026) 'DOE adds 12 Phase II Genesis Mission projects and six more Phase I awards', *R&D World*, 13 October. Available at: https://www.rdworldonline.com/doe-adds-12-phase-ii-genesis-mission-projects-and-six-more-phase-i-awards/

340. APLU (2026) 'Resources on APLU-AAU-ACE legal action contesting cuts to F&A reimbursement rates', updated 7 April. Available at: https://www.aplu.org/our-work/4-policy-and-advocacy/research-and-science/resources-on-aplu-aau-ace-legal-action-contesting-cuts-to-fa-reimbursement-rates/

341. Congressional Research Service (2026) 'NIH indirect costs policy for research grants: recent developments' (IN12516), 17 April. Available at: https://www.everycrsreport.com/files/2026-04-17_IN12516_be9b78de7dbbb7220ddf69779afc193119178fe5.html

342. Council of Graduate Schools (2026) 'Understanding the OMB's proposed rule on guidance for federal financial assistance', 15 June. Available at: https://cgsnet.org/insight/understanding-the-ombs-proposed-rule-on-guidance-for-federal-financial-assistance

343. FABBS (2026) 'President's FY27 budget request slashes funding for federal science agencies, again', 9 April. Available at: https://fabbs.org/news/2026/04/fy27-presidents-budget-request-slashes-funding-for-federal-science-agencies/

344. AAMC (2026) 'House subcommittee approves FY27 NSF spending bill', 1 May. Available at: https://www.aamc.org/advocacy-policy/washington-highlights/house-subcommittee-approves-fy27-nsf-spending-bill

345. Villafranca, M. (2026) 'FY2027 appropriations update: DOE Office of Science and NIEHS', AGU, 15 June. Available at: https://thebridge.agu.org/2026/06/15/fy2027-appropriations-update-doe-office-of-science-and-niehs/

346. National Institutes of Health (2026) 'NOT-OD-26-131: NIH operates under a continuing resolution', 28 September. Available at: https://grants.nih.gov/grants/guide/notice-files/NOT-OD-26-131.html

347. Department for Science, Innovation and Technology (2026) *Digital and Technologies Sector Plan: Year One Update*, June. Available at: https://assets.publishing.service.gov.uk/media/6a295455db0592d809052c11/D_T_Sector_Plan_PDF_F_A.pdf

348. Research Ireland (2026) *Curiosity, Capability, Competitiveness: Charting Ireland's Research and Innovation Future 2026–2030* (strategy snapshot), March. Available at: https://media.researchireland.ie/wp-content/uploads/2026/03/Summary-Strategy.pdf

349. Department of the Taoiseach (2026) 'Government publishes new Digital and AI Strategy', press release, 18 February. Available at: https://www.gov.ie/en/department-of-the-taoiseach/press-releases/government-publishes-new-digital-ai-strategy-90-actions-to-strengthen-irelands-position-as-a-digital-leader-and-ai-hub/

350. Hogan, D. (2024) 'Denmark launches leading sovereign AI supercomputer to solve scientific challenges with social impact', NVIDIA, 23 October. Available at: https://blogs.nvidia.com/blog/denmark-sovereign-ai-supercomputer/

351. Boudreau, K.J., Guinan, E.C., Lakhani, K.R. and Riedl, C. (2016) 'Looking across and looking beyond the knowledge frontier: intellectual distance, novelty, and resource allocation in science', *Management Science*, 62(10), 2765–2783. doi:10.1287/mnsc.2015.2285. Available at: https://ideas.repec.org/a/inm/ormnsc/v62y2016i10p2765-2783.html

352. Times Higher Education (2025) 'Reeves uplifts Aria budget to £1 billion and funds AI courses', June. Available at: https://www.timeshighereducation.com/news/reeves-uplifts-aria-budget-ps1-billion-and-funds-ai-courses

353. Advanced Research and Invention Agency (2026) 'How we work'. Available at: https://aria.org.uk/how-we-work

354. European Research Council (2026) 'ERC at a glance'. Available at: https://erc.europa.eu/about-erc/erc-glance

355. European Commission (2026) 'EIC Pathfinder', European Innovation Council. Available at: https://eic.ec.europa.eu/eic-funding-opportunities/eic-pathfinder_en

356. Research Ireland (2026) 'Research Ireland Investigators Programme'. Available at: https://www.researchireland.ie/funding/investigators/

357. Central Statistics Office (2026) 'Data for researchers', 12 August. Available at: https://www.cso.ie/en/aboutus/lgdp/csodatapolicies/dataforresearchers/

358. US Census Bureau (2026) 'Federal Statistical Research Data Centers', revised 9 September. Available at: https://www.census.gov/about/adrm/fsrdc.html

359. European Commission (2025) 'Commission adopts delegated act on data access under the Digital Services Act', 2 July. Available at: https://digital-strategy.ec.europa.eu/en/news/commission-adopts-delegated-act-data-access-under-digital-services-act

360. HM Revenue & Customs (2025) *HMRC Research Future Strategy*, August. Available at: https://www.gov.uk/government/publications/hmrc-research-future-strategy/hmrc-research-future-strategy

361. Irish Post (2026) 'Ireland's largest research institute will undergo major €100m expansion', January. Available at: https://www.irishpost.com/news/irelands-largest-research-institute-will-undergo-major-e100m-expansion-303168

362. Chapman, I. (2026) 'Ian Chapman maps the changes to UKRI's budget between 2026 and 2030', *Wonkhe*, 3 March. Available at: https://wonkhe.com/blogs/ian-chapman-maps-the-changes-to-ukris-budget-between-2026-and-2030/

363. UK Research and Innovation (2026) *UK Research and Innovation Strategy 2026 to 2031*, 13 July. Available at: https://www.ukri.org/publications/uk-research-and-innovation-strategy-2026-to-2031/

364. Taylor, J. (2025) 'Ruse or reform? Scientists debate Trump's "Gold Standard Science" order', *FYI*, American Institute of Physics, 1 July. Available at: https://www.aip.org/fyi/ruse-or-reform-scientists-debate-trumps-gold-standard-science-order

365. UK Research and Innovation (2024) 'Use of generative artificial intelligence in application preparation and assessment'. Available at: https://www.ukri.org/publications/generative-artificial-intelligence-in-application-and-assessment-policy/use-of-generative-artificial-intelligence-in-application-preparation-and-assessment/

366. Resultsense (2026) 'UK research funders to allow AI in grant assessment', 19 June. Available at: https://www.resultsense.com/news/2026-06-19-ukri-wellcome-ai-grant-assessment/