SURE VALLEY VENTURES · BRIEFING

AI at the Core

A briefing for university leaders

Barry Downes, Managing Partner, Sure Valley Ventures

RELEASE VERSION 1.0

24 November 2026

A companion to my full report, "AI at the Core". Two further companions, Beyond Adoption and AI at the Core: Who Checks the AI?, will follow. Footnotes give the main sources. Inside: what changed (p. 2) · the stack (p. 4) · costs (p. 7) · asks of governments (p. 9). Comments are welcome to barry@svv.ai.

Executive summary

AI is becoming the core method of science. In mathematics, its claimed results now arrive faster than experts can check them. On 26 October, OpenAI and Anthropic each claimed, independently, that its AI had found an algorithm that would refute the Exponential Time Hypothesis. Hundreds of results in theoretical computer science rest on that conjecture. Both proofs pass checks in Lean, a system that verifies proofs by computer. By Anthropic's own account, however, no person had checked its Lean code when it was announced. Volunteers have since rerun both proofs and read their formal statements, and an independent review of OpenAI's proof is under way. Other results have already passed review. In May, leading mathematicians verified an AI disproof of the Erdős unit distance conjecture. On 20 November, a proof of the Riemann hypothesis, found with Claude and written in Lean, passed independent machine checks within a day, and the next day the Lean and Mathlib teams 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.

Within five years, I expect AI agents, directed by small teams of researchers, to do most of the routine computational work of research. The change will be broad, reaching from mathematics and computer science through genomics, materials and engineering simulation, to economics, the social sciences and the humanities. Within ten, wherever experiments can be automated and results checked, the loop from hypothesis to validated result is likely to run with little human intervention. No university needs to train frontier models, but each needs access to AI compute, whether owned or shared.

The larger prize is science that could not be attempted before: questions too large, slow or uncertain to pursue, and connections across disciplines and facilities that no single group could make. Advantage will go to universities that own what sets them apart: distinctive data, automated experimentation, verification and the people to direct them. With these, they could advance discovery faster, more rigorously and with more real-world impact than larger institutions without them. In practice, that means three things: a right-sized owned or shared node alongside national and international machines; one institution-wide agreement for frontier models that protects data and intellectual property; and an accountable owner with authority to pause AI work if safety or security requires it. Running costs, mostly staff, can match the capital cost of a small stack every year, so they must be planned from the start.

For governments, the test of AI investment is the new science, companies and economic value it produces, more than the GPU hours it buys. Capital programmes buy equipment, but few pay to run university AI infrastructure beyond a project's life. Governments and their agencies should fund and build discovery pipelines around national facilities and data, with universities as nodes. They should also fund university operations and refresh, not only capital, and treat AI as a method for all of science, with a re-

sponsive route for radical ideas from anyone. Funders should set expectations for verifying AI-generated research, and universities should build the capability to meet them. What is built should be easy for start-ups and SMEs to use.

Adoption buys parity, not advantage. The opportunity is to attempt what was previously out of reach.

What changed in 2026

The frontier AI companies are building their own laboratories, so that experiments their models design can be run and the results fed straight back into the models. Anthropic now operates a wet biology laboratory, and OpenAI has run closed-loop experiments in an automated cloud laboratory.[1] The US, the UK, the EU and China have all made AI for science a national priority.[2] Researchers, meanwhile, are adopting AI faster than their institutions are supporting it.[3] Mathematics moved first because its results can be checked by machine. Biology and physics are following more slowly, because their results need laboratory validation.

When (2026)	Result	Status
January	Google DeepMind's AlphaGenome predicts the effects of genetic variants; predictions for about 9 billion variants later released for research use[4]	Peer reviewed (Nature)
May	An OpenAI model disproves the Erdős unit distance conjecture[5]	Verified: leading mathematicians published a checked version
September	Anthropic's Claude agents produce the first complete machine-checked proof of Fermat's Last Theorem, in about 11 days[6]	Machine-checked in Lean; formalises existing mathematics rather than proving new results
September	OpenAI says its agents proved finite-time blow-up for the Navier–Stokes equations, which it says meets the Clay Millennium Prize conditions[7]	Claimed; not peer reviewed; Clay lists the problem as active and has not awarded the prize; credit for the result disputed
October	Pinal, from a team led by Westlake University, designs proteins from plain-language descriptions of their function[8]	Peer reviewed (Nature); intended activity in all four classes tested; most screening computational
October	OpenAI and Anthropic each claim, independently, to have refuted the Exponential Time Hypothesis[9]	Machine-checked in Lean; volunteers have rerun both proofs and read the statements; independent review of OpenAI's proof under way
November	AutoDNA, an AI-native laboratory at Shanghai Jiao Tong University, runs more than 11,000 experimental steps across more than 20 types of instrument, with people only maintaining the equipment[10]	Peer reviewed (Nature)

1. Bort, J. (2026) 'Anthropic says its biology lab has already found something big', *TechCrunch*; OpenAI (2026) 'GPT-5 lowers the cost of cell-free protein synthesis'.
2. The White House (2025) *Executive Order 14363: Launching the Genesis Mission*; Department for Science, Innovation and Technology (2025) *AI for Science Strategy*; European Commission (2025) *European Strategy for AI in Science*; State Council of the People's Republic of China (2025) 'Opinions of the State Council on deepening the implementation of the "Artificial Intelligence+" initiative', translated by CSET.
3. Wiley (2025) 'AI adoption jumps to 84% among researchers as expectations undergo significant reality check'; Springer Nature (2026) 'Early career researchers lead shift towards AI-first discovery'.
4. Avsec, Ž. et al. (2026) 'Advancing regulatory variant effect prediction with AlphaGenome', *Nature*, 649, 1206–1218. doi:10.1038/s41586-025-10014-0.
5. OpenAI (2026) 'An OpenAI model has disproved a central conjecture in discrete geometry'; 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.
6. Anthropic (2026) 'Formalizing Fermat's Last Theorem'.
7. OpenAI (2026) 'On the Navier–Stokes Millennium Prize Problem'; Clay Mathematics Institute (2026) 'Navier–Stokes equation', Millennium Prize Problems; Kakaes, K. (2026) 'AI has solved one of math's $1 million Millennium Prize Problems', *Quanta Magazine*.
8. Dai, F. et al. (2026) 'Programmable design of functional proteins from natural language', *Nature*.
9. OpenAI (2026) 'Working with academic communities to build scientific understanding'; Anthropic (2026) 'A 2^o(n) time algorithm for 3-SAT' [GitHub repository: anthropics/formal-math, eth/README.md]; Impagliazzo, R., Paturi, R. and Williams, R. (2026) [Open letter on the independent review of the claimed refutations of the Exponential Time Hypothesis], October.

Taken together, these results show that fields where answers can be checked quickly move first, and that claims can run ahead of evidence, which gives universities a natural role as independent verifiers. 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.[11] Model access is opening up through free and discounted academic programmes, although the most capable models are often held back, especially in the life sciences[12] (full report, Sections 2 and 3).

When I say that "AI is becoming the core method of science", I mean a general-purpose instrument used across the research cycle, as statistics and computing became, not a replacement for experiment or interpretation. By AI, I mean the whole family, from large language models (LLMs) to domain models such as AlphaFold and world models, which learn how the physical world behaves. The same capabilities raise real risks, whose severity experts dispute,[13] and that is one reason to build verification, governance and security research into the infrastructure.

What is changing, field by field

AI moves fastest where there is good data, a quick way to check answers and a way to act on predictions, so progress is uneven across fields. In all of them, it works best as an instrument checked against data and expert judgement.

Field	Where AI stands in late 2026	Where universities are needed
Life sciences and medicine	Designing proteins and antibodies; genome models; AI-designed drugs in late-stage trials; automated laboratories	Wet-laboratory validation, cell-response data, biosecurity
Mathematics	Machine-checked proofs; new constructions and algorithms; open problems solved	Expert judgement of what formal results mean
Chemistry, materials and engineering	Generative design and robotic synthesis; AI surrogates for engineering simulation; some high-profile AI claims corrected after independent checks	Synthesis, characterisation, testing, proving novelty
Physics, energy and quantum	Fusion plasma control, particle identification, quantum error correction and circuit design	Access to facilities and hardware
Earth, climate and agriculture	AI weather forecasts in operational use; climate emulators and precision agriculture emerging	Extrapolating beyond observed conditions
Social sciences and humanities	Reading damaged ancient texts; evidence synthesis; simulated respondents	Validity against real human data
Computer science and AI	Coding agents write much new code; alignment, interpretability and security are open research fields	Compute and model access; verifying code written by agents

The life sciences offer the largest gains. AI-engineered antibodies are moving into late-stage trials,[14] genome models trained on trillions of DNA letters can predict the effects of variants without task-specific training,[15] and AI-supported breast screening detected 29% more cancers in a randomised trial in Sweden.[16] Models of how cells respond to change do not yet consistently beat simple baselines, largely for lack of data.[17] In October 2026, the US Department of Energy, Biohub, Google DeepMind, Isomorphic

10. Wu, M. et al. (2026) 'An AI-native experimental laboratory for autonomous biomolecular operations', *Nature*.
11. METR (2026) 'Independent assessment of AI R&D acceleration at Anthropic in October 2026'.
12. Anthropic (2026) *Expanding support for scientists*; OpenAI (2026) 'Accelerating scientific discovery with ChatGPT for Academic Researchers'; Axios (2026) 'OpenAI launches new AI model for life sciences research'.
13. Bengio, Y. et al. (2026) *International AI Safety Report 2026*.
14. 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'.
15. Brixi, G. et al. (2026) 'Genome modelling and design across all domains of life with Evo 2', *Nature*, 652, 1349–1361.
16. Lund University (2025) 'AI-supported breast cancer screening – new results suggest even higher accuracy' (MASAI trial, *The Lancet Digital Health*).

Labs and Meta committed about $1.3 billion of new money to AI-ready biological data, with the National Institutes of Health coordinating existing datasets (nearly $1.8 billion in all).[18] The life sciences also carry the most serious risks. Providers already restrict their most capable models for biology, so universities will need controls such as vetting users, screening orders for synthetic DNA and reviewing dual-use work.[19]

What research may look like in 2031 and 2036

By 2031, I expect a typical research group 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. Grants will budget compute like equipment, and compute per researcher will become a recruitment factor.

By 2036, research capacity will depend on infrastructure as much as headcount. Universities will be valued for what is hard to replicate: distinctive instruments, field sites, patient cohorts and regional industry data; trusted verification; and the education of people able to direct machine research. National facilities, from synchrotrons to health, Earth-observation and climate data, are likely to become discovery engines that universities, public programmes and companies can join at the right layer rather than each starting from scratch.

These forecasts are uncertain. General-purpose laboratory robotics may lag (even though specialised systems such as the AutoDNA laboratory show what is already possible), specialist scientific models may not generalise beyond their training data, compute and energy may stay scarce, the next wave of AI may not run on today's models or hardware, private laboratories with far more capital and compute may outrun universities, and gains in output may prove smaller than gains in adoption. I do not think any of these would reverse the direction of travel, but each could slow it (full report, Section 5).

Build a stack, own what differentiates

A university and its innovation campus should build an AI research stack in layers. It should own the layers that give lasting advantage and rent or federate the rest.

17. Arc Institute (2025) 'Virtual Cell Challenge 2025 wrap-up: winners and reflections'.
18. 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'.
19. Anthropic (2025) 'Activating ASL-3 protections'; Axios (2026) 'OpenAI launches new AI model for life sciences research'.

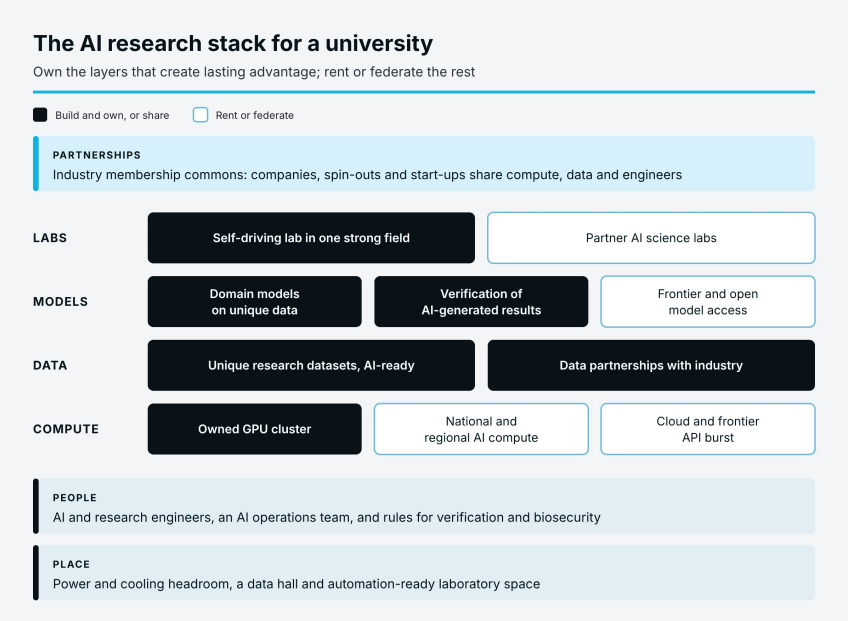

Figure. The AI research stack. The dark layers are the ones that give lasting advantage, owned or held with a partner.

- **Compute in three tiers.** First, where measured workloads, sensitive data or services to industry justify it, an owned cluster for inference, fine-tuning, teaching and sensitive data, sized to that demand and planned to complement national systems. Princeton funded a 300-GPU cluster largely from its endowment.[20] A regional university can start with 16 to 32 GPUs and grow towards about 64. Second, national and international machines for large runs. Third, one institution-wide agreement for frontier models, rather than many individual ones, covering confidentiality, no training on users' work, ownership of outputs and research priority. Most work can run under the education and academic agreements that Google, Microsoft, OpenAI and Anthropic already offer.[21] Keep open models on the owned cluster for patient, partner-confidential and export-controlled data, and for continuity if a provider's models are withdrawn, and negotiate text-and-data-mining rights.
- **Models as a shared service.** A curated catalogue of open models and a small central team to adapt, deploy and evaluate them, with domain models in the two or three fields where the institution has distinctive data.
- **Data as the differentiator.** Access to capable general models is widely shared; distinctive data is not. Inventory what the university and its region uniquely hold, prepare it for AI use under proper governance, and build data partnerships with industry.
- **A closed-loop laboratory.** At least one self-driving laboratory, where AI designs experiments, robots run them and results inform the next experiments. Liverpool, Illinois and Toronto show that universities can build these,[22] and low-cost platforms now make a pilot affordable.
- **Verification and governance.** Machine-checked results where possible, named expert review of what formal results claim, auditable records of agent work, and controls for provenance and biosecurity. Governance spans every layer and works best as one function. That function needs an accountable owner and a small group with authority to pause AI work when, for example, an agent acts beyond its permissions, sensitive data is exposed or a result raises dual-use concerns. Its remit covers agent permissions, ethics review, data protection, EU AI Act duties, dual-use screening and compute allocation.

20. Princeton University (2024) 'Princeton invests in new 300-GPU cluster for academic AI research'.
21. Google (2025) 'Gemini app now available with enterprise-grade data protections for all Google Workspace for Education users'; University of Cambridge Information Services (2024) 'Copilot', IT Help and Support.
22. Dai, T. et al. (2024) 'Autonomous mobile robots for exploratory synthetic chemistry', *Nature*, 635, 890–897; University of Illinois News Bureau (2025) 'Self-driving lab: AI and automated biology combine to improve enzymes'; University of Toronto Engineering News (2023) 'Self-driving labs: $200-million federal grant powers AI-driven materials discovery…'.

- **People.** A career track for research software and AI engineers, with standing comparable to academic posts, an AI operations function, and doctoral training that teaches students to direct and check AI without losing the underlying skills.
- **Place.** Power and cooling headroom well beyond the first node (megawatts if the campus will host industry or national capacity), a data hall, and laboratory space designed for automation, planned with the institution's net-zero commitments in mind.

The emphasis differs by institution (full report, Section 6.2).

- **Research-intensive universities** should prioritise an owned cluster of meaningful scale, domain models in their strongest fields and a leading role in national programmes.
- **Regional and technical universities** should lead with distinctive regional and industry data, applied domain models, an operations team built from existing engineering staff, and a right-sized owned node, or a share in one hosted by a partner institution that already has the facilities to run it. Their strength is close relationships with industry, and the main risk is spreading resources too thin.
- **Private and endowed universities** can fund a cluster from endowment, but must plan for running costs and limits on private use of bond-financed facilities.
- **Innovation campuses** can make the stack a central part of their offer, and a good test is whether companies choose the campus for what it offers, not only for its buildings.

Seven questions for 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?

Intellectual property and licensing

Technology transfer offices will face three questions first (full report, Section 6.11).

- **Who is the inventor?** Courts and patent offices in the UK, Europe and the US agree that an inventor must be a natural person,[23] and US guidance issued in November 2025 treats AI systems as tools used by human inventors.[24] Record who framed the problem, directed the agents and recognised the result, as evidence of who conceived the invention.
- **How should models and data be licensed?** First check that the university has the right to license: health-service data, partner-owned data and journals' data requirements often limit it. Code, data and model weights can then each carry their own licence, from standard open source to open li-

23. UK Supreme Court (2023) *Thaler v Comptroller-General of Patents, Designs and Trade Marks* [2023] UKSC 49; European Patent Office (2021) 'Legal Board of Appeal decision J 8/20 (DABUS)'; US Court of Appeals for the Federal Circuit (2022) *Thaler v. Vidal*, No. 2021-2347, 43 F.4th 1207.
24. US Patent and Trademark Office (2025) 'Revised inventorship guidance for AI-assisted inventions', *Federal Register*, 90 FR 54636.

cences with use restrictions or free research use with commercial use licensed separately. Under the EU AI Act, expect developers to ask about the provenance of any dataset you license to them.[25]

- **How is IP handled in a shared facility?** Models range from the Hartree Centre's template, in which each party keeps its background IP and owns what its people create,[26] to the Structural Genomics Consortium, which never patents.[27] A membership commons needs these terms agreed before the first member joins.

Standard agreements have not kept pace, and 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.[28] Until guidance catches up, clauses will need to be drafted.

What it costs

Capital runs from about €0.5M to €0.9M for a 16-GPU node to about €14.4M ($16M) for a 256-GPU cluster, mostly for the GPUs themselves.[29] A minimal 16-GPU node, run by two or three engineers, costs roughly €0.25M to €0.35M a year. The full stack in the table costs €1.3M to €3.2M a year, before maintenance contracts and model access. Staff are the largest recurring cost. The staff and power figures below are my estimates, at about €100,000 a year fully loaded per person and about €0.20 per kWh (€1 ≈ $1.11 ≈ £0.84; full report, Section 7).

Layer	Typical starting scale	Capital	Annual running cost
Owned GPU cluster	16 to 256 Hopper-generation NVIDIA GPUs (H100/H200)	€0.5M–€14.4M	Maintenance contracts
Power and cooling	2 to 32 eight-GPU systems at up to about 10 kW each	–	€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	–	€0.2M–€0.4M
Verification and research integrity	2 to 3 specialists	–	€0.2M–€0.3M
AI governance, risk and compliance	1 to 2 specialists supporting an existing senior owner	–	€0.1M–€0.2M
Self-driving laboratory in one domain	From open-hardware rigs to a research-grade robot scientist	€0.05M–€1.2M	€0.1M–€0.3M
Frontier model access	Institution-wide	–	Varies; academic programmes exist

The capital figures cover hardware only; the running costs assume a new team. Extending an existing data centre costs far less at the outset. Plans should show both the incremental cost and the full economic cost over five years, including replacing shared plant such as cooling.

25. European Commission (2025) 'The General-Purpose AI Code of Practice'.
26. STFC Hartree Centre (2023) *HNCDI EXCELERATE IP Template Agreement*, June.
27. Structural Genomics Consortium (n.d.) 'About'.
28. Knowledge Transfer Ireland (2019) *National IP Protocol*; Intellectual Property Office (2022) 'Lambert Toolkit'.
29. IntuitionLabs (2026) 'Data center GPU pricing 2026: the full AI pricing index'.

Running costs are the hardest part to fund. Capital grants rarely pay to operate the equipment they buy (full report, Section 7.4).

- Ireland's €750M INSPIRE programme does not fund operating salaries or maintenance beyond a project's life. Research Ireland's 2025 infrastructure programme funds two years of maintenance and expects access charges to cover running costs.[30]
- The EuroHPC Joint Undertaking can meet up to 50% of the operating costs of the high-end supercomputers it co-owns, and 35% for mid-range systems such as Ireland's forthcoming CASPIr.[31]
- The UK funds up to £2M a year to run each of a few selected National Compute Resources, but nothing equivalent for other university systems.[32]
- In the US, NSF funds operating costs of up to 15% of the acquisition cost a year, but only for a handful of selected national resources.[33]

University-hosted infrastructure outside those systems generally has no equivalent. Running costs must therefore come from industry membership, grant overheads, research centre roles and paid access, and they need planning before any equipment is bought.

Sharing with industry has limits. Under EU State aid rules, commercial use that relies on the same inputs as the research and takes no more than 20% of a facility's annual capacity can remain ancillary and outside State aid control.[34] Research Ireland builds this into its infrastructure grants, which gives grantees a clear framework: paid commercial use of up to 20% of annual capacity, at full market rate, while qualifying collaborative research with companies counts as academic use.[35] For AI compute, 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 a sustained commercial service (full report, Section 7.5).

The UK's subsidy control regime asks instead whether commercial use is limited in scope, giving less than 20% of capacity as an example rather than a rule.[36] In the US, federally funded equipment may not undercut private providers and bond-financed facilities face caps on private use.[37] Larger commercial use usually needs separately funded capacity, which is a natural opening for co-investment by developers, investors and industry partners.

Access alone is rarely enough: companies, and SMEs in particular, want expertise. Offer collaborative research, testing and specialist services, with people whose job is to serve companies, as Ireland's Technology Gateways and the UK's Knowledge Transfer Partnerships do.[38]

Build in stages, and plan for refresh. Start by making the AI compute that already exists shareable, then add capacity as demand grows. Public capital arrives through calls every year or two, so plan each call as the next stage of one programme. Later tranches will buy better and more varied hardware for the money than today's shortage prices allow, as alternatives to NVIDIA, such as AMD's accelerators and Google's TPUs, mature. Older GPUs remain useful for inference, fine-tuning and teaching long after they

30. Higher Education Authority and Research Ireland (2026) *INSPIRE Research Infrastructure Programme Phase One: Call Document*, May; Research Ireland (2025) *Research Ireland Infrastructure Programme 2025: Call for Submission of Proposals*.
31. Council of the European Union (2021) 'Council Regulation (EU) 2021/1173 establishing the European High Performance Computing Joint Undertaking', Articles 11(2) and 14(2) (since amended).
32. Department for Science, Innovation and Technology (2025) *UK Compute Roadmap*; UKRI (2025) 'Phase 2 National Compute Resource: expression of interest'.
33. National Science Foundation (2024) *Advanced Computing Systems & Services: Adapting to the Rapid Evolution of Science and Engineering Research* (solicitation NSF 24-583).
34. 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, paragraph 21.
35. Research Ireland (2025) *Research Ireland Infrastructure Programme 2025: Call for Submission of Proposals*.
36. Department for Business and Trade (2025) *UK subsidy control regime: statutory guidance*.
37. US Code of Federal Regulations, 2 CFR § 200.313, 'Equipment'; 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.
38. Enterprise Ireland Technology Gateway Network (2025) 'What are Gateways?'; UK Research and Innovation (2025) 'Knowledge Transfer Partnership guidance: who can take part', Innovate UK.

leave the frontier.[39] Keep software portable, and put the larger, longer-lived investment into data, people and the new verification capabilities that AI-generated research requires.

A three-year start for a regional university

1. **Year 1: foundations.** Appoint one programme owner. Take stock of compute and data, including capacity that could be shared, choose two or three priority domains and secure national and international access. Build a small node sized to measured demand, or share one hosted by a partner institution. Hire the first four to six of an eventual six to 12 research engineers. Name an accountable owner for AI governance, with authority to pause AI work that breaches policy or creates unacceptable risk. Adopt policies on research data, ethics, research security and the use of AI agents.
2. **Year 2: first assets.** Prepare one distinctive dataset with an industry partner; adapt an open domain model on it; pilot a low-cost self-driving laboratory; launch an industry membership offer for the campus.
3. **Year 3: scale on evidence.** Grow the cluster in tranches, towards about 64 GPUs if use justifies it; move to a research-grade laboratory only if the pilot has produced data and partners; bid to host a regional node in national programmes, using the track record built in the first two years.

Plan the stack as one programme, even though it will usually be funded through a series of bids, and match each funding instrument to the layer it suits. Funders score the capacity to deliver.

What to ask of governments

The US, the UK and the EU, with Ireland through EU programmes and its own investments, agree on the direction. In Ireland, AI research infrastructure could help deliver Research Ireland's 2026–2030 strategy, which targets 14 research centres and 50 active spin-out companies.[40] National platforms, however, are still owned and governed centrally, and outside a few nationally selected systems none systematically funds the running costs of university AI infrastructure. The full report proposes six changes (Section 8).

1. **Treat AI as a method for all of science**, opening national programmes to every scientific discipline, from mathematics and physics to economics and the humanities, with a strand for high-risk work and a responsive route for radical ideas from anyone, which I propose should take about a quarter of new public funding for AI-for-science programmes, including the compute and data access they provide.
2. **Build AI discovery pipelines around national facilities and data**, such as synchrotrons, health and climate data, and new, securely held economic and social data, with universities, including regional ones, as funded nodes and routes in for start-ups.
3. **Fund university operations and refresh, not only capital**, with a recurrent operations budget tied to measured use and capital calls planned as stages of one programme, and publish what national compute achieves, not just GPU hours.
4. **Set expectations for verifying AI-generated research** that combine machine checking with named expert review, with governance standards and model contract clauses for AI providers.
5. **Secure sector-wide access to frontier models on research terms**, covering price and protection of data and IP, governed access for life-sciences researchers and access for university alignment, evalu-

39. Koundinya, S. (2026) 'Why NVIDIA's six-year-old GPU is still making money', *Analytics India Magazine*.
40. Research Ireland (2026) *Curiosity, Capability, Competitiveness: Charting Ireland's Research and Innovation Future 2026–2030* (strategy snapshot), March.

ation and security research, alongside open-weight models and nationally controlled capacity, so research can continue if access is restricted.

6. **Make it simple for companies to share university infrastructure**, including, in Ireland, funding separate capacity for industry or reviewing how the 80:20 commercial-use condition applies to AI compute, and fund research engineering careers and AI-fluent doctoral training at scale.

Four of the six can be delivered through instruments that already exist:

Change	Ireland	United Kingdom	European Union	United States
Discovery pipelines and university nodes	Nodes linked to CASPIr and the AI Factory Antenna; AI-ready CSO data	Pipelines around facilities such as Diamond; regional AIRR nodes; AI-ready economic and social data	University nodes in the RAISE network; platform data under the DSA	Regional university nodes in NAIRR and Genesis
Fund operations and refresh	An operations strand alongside INSPIRE capital	Extend the National Compute Resource model to more, and smaller, university AI nodes	Operating support for university nodes through FP10 and the EuroHPC Joint Undertaking	Wider NSF-style operations allowances
Verification	Research Ireland, HEA and National Research Integrity Forum expectations	UKRI research-integrity expectations	ERA Living Guidelines in FP10 grant conditions	NSF and NIH grant conditions
Frontier model access	A national academic agreement	A sector-wide agreement	An agreement through RAISE	Academic programmes on agreed terms

CSO: Central Statistics Office; HEA: Higher Education Authority; AIRR/NAIRR: UK/US national AI research resources; RAISE: EU AI-for-science resource; FP10: next EU framework programme; DSA: Digital Services Act; CASPIr: Ireland's forthcoming EuroHPC system; ERA: European Research Area.

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.

Interests. This briefing 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 briefing. 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 the full report. SVV is an investor in Astut Ltd, an Oxford spin-out of which Professor Pete Grindrod is Chief Executive. This briefing 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.