UNIVERSITY OF CAMBRIDGE

ai@cam

ai.cam.ac.uk

From access to adoption: case studies of AI adoption in science institutions

In partnership with:

Bennett School of Public Policy Cambridge

CONTENTS

Introduction	**4**
Themes emerging from the case studies	**6**
Building the missing middle in AI for science	**8**
Case Studies	**9 - 37**
▸ The University of Cambridge's Accelerate Science initiative	**10**
▸ KU Leuven's Leuven.AI programme	**13**
▸ The University of Oxford's AI in Science fellowships	**16**
▸ Data Science Africa (DSA)	**19**
▸ The Machine Learning for Science Cluster at the University of Tübingen	**22**
▸ The Alchemy hub for AI in chemistry	**25**
▸ The UK Centre for Ecology and Hydrology (UKCEH)	**28**
▸ The European Centre for Medium-Range Weather Forecasts (ECMWF)	**31**
▸ The OpenBind consortium	**34**
▸ Polymathic AI	**37**
Contact	**40**

The UK starts from a strong position in AI for science, with world-leading scientific institutions, facilities, and datasets alongside substantial AI research capability and public investment in compute infrastructure. Policy interventions to date have largely focused on strengthening these inputs. Their scientific value will also depend on how effectively new capabilities diffuse into research practice.

This brief examines ten initiatives that are approaching this challenge in different ways. Across them, adoption relies on institutional capabilities that receive less attention in the AI for science policy debate: engineering and data stewardship; routes for researchers to experiment with new tools; mechanisms for testing and assurance; communities through which expertise spreads; and support for successful experiments to become sustained services and infrastructure.

Adoption requires both technical and institutional innovation, as institutions develop new roles, services, and working practices around AI, discover where it is useful, and find technical limitations that drive further innovation. The UK's scientific institutions are important sites for this process.

As AI for science policy moves from strategy to implementation, the case studies described here point to questions that will shape whether the UK's investments translate into scientific and public value:



When government invests in a national capability, what provides the institutional machinery that makes it usable across the whole research base?



How can public funding sustain the capacity to learn in a field changing this quickly?



What allows a successful innovation – institutional or technical – to scale, and who is responsible for it once it has?



How can the UK ensure that AI delivers public value after decades of public investment in scientific resources?

Introduction

Science has provided some of the highest-profile demonstrations of AI's progress, from protein structure prediction to weather forecasting.

As AI capabilities become more widely accessible, AI for science policy questions are shifting from what AI can do for science to how science changes as researchers begin to use it.

This involves changes in research practice, technical roles, scientific workflows, and, in some cases, the institutions through which research is organised. There are growing signs of disruption across areas of research. Mathematicians are debating the future of their discipline;[1] arXiv and other publishers are rewriting their paper submission policies;[2] and there are indications that AI use is influencing public confidence in science.[3] While scientists report substantial time savings from AI use, these come with growing backlogs of untested hypotheses and increased demands for verification, suggesting that individual gains might lead to bottlenecks elsewhere in the research process.[4]

The prospect of increasingly autonomous science intensifies these questions. AI systems can formulate hypotheses, design experiments, analyse results, and coordinate research workflows, prompting predictions that substantial parts of scientific discovery could be automated. Today's agentic AI systems[5] point towards the technical foundations for this future, while exposing the work required to make scientific data and models usable across disciplines. Advances also depend on connecting new computational capabilities with scientific knowledge and judgement.[6] Science has a subjective element,[7] with researchers making choices about which questions are worth pursuing, what constitutes sufficient evidence, and how findings should be interpreted and used. How these choices are made matters for the credibility of science, and for public confidence in its findings. Whatever the eventual capabilities of automated science, the institutions through which research is conducted will shape what those capabilities are used for and whose interests they serve.

For policymakers, the hope has been that AI would accelerate breakthrough discoveries that drive economic growth and create public value, an ambition set out in the UK's AI for Science Strategy.[8] While AI policy has tended to focus on the things researchers need access to (data, compute, models, and skills), less attention has been paid to what happens when those resources meet research practice.

UK policy in AI for science to date has been weighted towards supply-side assets. The AI Research Resource (AIRR) is planned to expand. UKRI established nine AI research hubs in 2024 to build capability in AI methods and their application across disciplines. The Sovereign AI Unit invests in UK AI companies and in strategic capabilities for science, including AI for science-focused calls. The AI for Science Strategy, published in November 2025, made commitments on funding, data access, and skills development. ARIA's Activation Partners programme aims to bring AI for science tools and expertise to the researchers it funds. Whether these investments translate into scientific progress depends on what happens where they meet research practice.

The October 2026 Budget reinforced the Government's commitment to AI infrastructure and sovereign capability, introducing new measures to support AI adoption across the economy and announcing a new AI Biosecurity Lab that will combine public investment with access to commercial AI models. These developments sharpen the question of how investment in AI capabilities translates into sustained scientific capability, and what institutional arrangements are needed to support that transition.

The ten case studies in this brief examine how research institutions are responding to this challenge at different points in the development and adoption of AI for science. Some focus on making existing AI capabilities accessible to researchers, through engineering services, training, and fellowships. Others develop new scientific AI capabilities, including models, datasets, and operational systems. Across these different approaches, the cases illustrate how scientific institutions contribute both to the diffusion of AI and to innovation in the technologies themselves. They include:

- **The University of Cambridge's Accelerate Science initiative:** a team of machine learning engineers supporting cross-institutional skills-building and project incubation.
- **KU Leuven's Leuven.AI programme:** an engineering service offering short-term expertise to domain scientists across the University.
- **The University of Oxford's AI in Science fellowships:** a postdoctoral and faculty fellowship programme that provides funding, mentoring, and engineering support to interdisciplinary cohorts.
- **Data Science Africa's (DSA's) capability-building activities:** a grassroots network that runs summer schools and fellowships and supports local chapter activities across over a dozen African countries.
- **The Machine Learning for Science Cluster at the University of Tübingen:** a 14-year, nationally-funded cluster of excellence that links machine learning with domain and ethics researchers.
- **The AIchemy hub for AI in chemistry:** an EPSRC-funded hub led by Imperial College London and the University of Liverpool that incubates a mix of AI for chemistry projects.
- **The UK Centre for Ecology and Hydrology (UKCEH):** a national environmental research institute managing decades of monitoring data, and integrating AI into biodiversity monitoring research.
- **The European Centre for Medium-Range Weather Forecasts (ECMWF):** an intergovernmental forecasting centre that

is using a machine learning model in operational service alongside its physics-based models.

- **The OpenBind consortium, coordinated by the Diamond Light Source:** a Sovereign AI Unit-backed programme to generate the world's largest open dataset of protein-ligand interactions to develop AI models for drug discovery.
- **Polymathic AI:** an international collaboration to build foundation models that learn across scientific domains.

The case studies draw on the experiences of initiatives working at different scales and stages of development. They are not intended as comparative evaluations of programme effectiveness, but as examples of how institutions are responding to the practical challenges of AI adoption.

Each connects to core areas of policy activity, set out in Table 1. Across these different approaches, the cases show how institutional capabilities shape both the development of AI for science and its adoption in research practice. This brief examines what this means for the next phase of AI for science policy.

Table 1: Underlying capabilities of the AI for science system



Scientific foundations

High-quality scientific data; capacity to generate new data; experimental facilities; domain knowledge; standards and metadata

Relevant use cases:

OpenBind, UKCEH, Polymathic



Technical foundations

Compute; models; software; storage; tooling; interfaces; engineering capacity

Relevant use cases:

Polymathic, Accelerate, ECMWF, Leuven



People & practice

AI skills; domain–AI translation; research software engineering; data stewardship; interdisciplinary working; scientific judgement

Relevant use cases:

Leuven, Accelerate, Oxford, ECMWF, Alchemy



Institutions & incentives

Funding that accommodates experimentation; career structures; team-science incentives; organisational capacity; routes from prototype to sustained use

Relevant use cases:

Tübingen, Oxford, Alchemy, ECMWF



Networks & system infrastructure

Communities of practice; shared services; national facilities; disciplinary networks; international collaboration; mechanisms for diffusion

Relevant use cases:

DSA, Alchemy, Accelerate, ECMWF, Tübingen



Stewardship & assurance

Scientific validation and reproducibility; responsible research practices; human oversight and judgement; monitoring and contestability; environmental and computational sustainability; maintenance and lifecycle responsibility; management of critical dependencies

Relevant use cases:

ECMWF, UKCEH, OpenBind, Tübingen

Themes emerging from the case studies

THE GAP BETWEEN ACCESS AND ADOPTION

Access to AI capability does not guarantee its use. For example:

- While the University of Cambridge had investment for the Dawn HPC facility, it was not accessible for scientists until Accelerate Science designed an on-ramp programme.
- UKCEH needed to implement a data stewardship effort before its significant data assets were amenable to AI analysis.
- KU Leuven's researchers can identify promising AI for science problems, but require the engineering capacity to pursue them.

Compute illustrates the distinction between access and use. Some researchers need access to modest amounts of compute, coupled with the support to use it. Others are developing frontier models that require access to very large allocations. KU Leuven finds that most of its projects run comfortably on standard university computing, with the Flemish supercomputer reserved for the minority that need it. In contrast, Polymathic AI needed to assemble resources from the Flatiron Institute, NVIDIA, and France's Jean Zay supercomputer. National compute investment speaks most directly to this second need, but in both cases what determines use is the people and processes that connect researchers to the facilities they need.

Much of the work between access and use is done by intermediaries: engineers, data stewards, and staff providing hands-on support. These roles can be organised in different ways. Accelerate Science and KU Leuven have developed shared engineering services to support researchers across their institutions, while ECMWF embeds machine learning expertise within scientific teams. In each case, dedicated technical capacity helps connect AI capabilities to scientific problems.

The research system does not have a settled place for many of the people doing this connective work. Machine learning engineers face limited academic career pathways and salary competition from industry, while engineering and data stewardship remain substantially less visible in academic reward structures than research publications. Those institutions located in areas with a high density of AI activities may be better placed to negotiate such challenges, creating a risk of inequality in adoption across institutions. The policy question that follows is, when government invests in a national capability, what provides the institutional machinery that makes it usable across the research base?

ADOPTION IS ALSO DISCOVERY

The question of 'where should the UK invest in AI for science?' is difficult to answer from the centre. Valuable applications are distributed across disciplines, while the technology is changing faster than conventional funding cycles. Researchers often know where the opportunities and obstacles are before funders do.

Experimentation makes some of this demand visible. Alchemy's call for funding in AI for chemistry was oversubscribed in its first run; DSA's summer school received 1,800 applicants for 400 places; Accelerate Science's courses are regularly oversubscribed. Adoption programmes therefore do more than support uptake: they signal where demand is emerging and which bottlenecks recur.

Successful programmes have also had space to learn as the scope of AI for science has changed. Accelerate Science has repeatedly updated its training offer as new AI capabilities appear, with modules on LLMs and agentic AI. Oxford University's fellowship programme developed an individual tuition approach. Tübingen's Colaboratory experimented with different delivery forms. Neither the most valuable applications nor the support they required could be fully specified in advance.

Such exploratory benefit therefore benefits from flexibility in funding. Alchemy left most of its £12m budget unallocated at launch, allowing it to fund ideas that did not exist when the hub was proposed. Tübingen's Cluster of Excellence combines a demanding selection process with seven-year funding periods and light oversight, so that researchers can adapt without justifying each change of direction. The question for policy is how public funding can preserve the flexible capacity to learn in a field changing this quickly.

SCALING SUCCESSFUL EXPERIMENTS

Many of the initiatives described here began as experiments: a new method, service, or way of working. The harder institutional problem begins when an experiment works. Who maintains it, pays its recurring costs, supports its users, and takes responsibility for it once researchers begin to depend on it? A grant can fund a new method or pilot service. If researchers begin to depend on it, someone must maintain the software, steward the data, support its users, and pay the recurring costs. Research funding is not generally designed for this transition.

ECMWF offers one example of this transition for a new AI for science tool. Its AI Forecasting System (AIFS) became operational after a period as an experimental model, which was run live but not used in operations, while national weather services tested it. This experimental phase allowed users to identify failure modes not captured by initial benchmark testing. Moving to operation required ECMWF to adapt its verification systems to catch AI failure modes and build user understanding of how to use the tool effectively.

Other forms of confidence-building or assurance can be seen across the case studies. OpenBind uses blind challenges to test model claims; UKCEH keeps a scientist in the loop in AI projects; Tübingen embeds ethics researchers in its AI teams. These projects are also finding new points of vulnerability as AI enters scientific workflows; for example, Polymathic found coding agents took shortcuts on scientific data.

Assurance is one part of the transition. Successful tools also need to be sustained. Tool maintenance, data stewardship, technical staff, and other recurring costs attract less funding and recognition than new methods. UKCEH describes how this transition falls between existing mechanisms: research grants fund new methods and government procures finished services, but the work of turning experimental tools into operational infrastructure does not fit easily in either.

The question for policy is what allows something that works to

become part of how science is done, even within the originating institution, and who is responsible for scaling it once it has succeeded.

SOVEREIGN ASSETS AND DEPENDENCIES

Our Navigating AI Sovereignty policy brief[9] examined the UK's sovereign AI position along the AI value chain (data, compute, models, talent, applications, and governance) and argued that sovereignty means having the ability to make choices across this stack, rather than owning every level. AI for science makes the sovereignty question more concrete. Policy interventions need to engage with where models or compute are located, and whether scientists retain the capacity to generate evidence, scrutinise AI outputs, and determine the questions their research pursues.

Many of the case studies here rely on scientific assets built through decades of public investment. OpenBind builds from the Diamond Light Source; UKCEH uses over 60 years of environmental monitoring data; the Aardvark project benefits from the UK Met Office's archive. These assets give UK researchers capabilities that cannot quickly be recreated elsewhere.

Despite decades of policy efforts to develop scientific data infrastructures, data remains a constraint on AI development. As the UKCEH and AIchemy projects show, preparing data for AI analysis is technically challenging, may not produce publications, and is hard to fund. Such curation does, however, have the potential for widespread impact: the Polymathic project describes how its data releases have been as catalytic for research as its model development.

The current dynamics of AI development create new tensions in the policy conversation about how to manage national data assets. Open release has been a default for publicly-funded scientific data, with the aim of returning benefits to research and to the public. AI complicates that assumption. If public datasets become strategic inputs to privately owned models, under what conditions does openness continue to maximise public value?

The cases show institutions managing these dependencies in various ways: ECMWF uses portable software across different compute environments; UKCEH combines external cloud services with cloud-independent development; and other initiatives use controlled archives, locally hosted models, or international compute partnerships. Sovereignty here is a question of preserving options rather than eliminating dependencies.

The cases also point to less visible dependencies in the scientific system. Data Science Africa's experience training researchers across the continent raises questions about where scientific capacity accumulates as researchers move into industry roles. Other cases illustrate dependencies arising from funding and institutional arrangements. Philanthropic support has enabled important UK AI for science initiatives, while public investments through ARIA and the Sovereign AI Fund have skewed towards international organisations. These arrangements create opportunities for scientific progress, but also raise questions about where capabilities accumulate, who determines research priorities, and how they can be sustained over time.

The policy question is how those making public investments should structure access to ensure a reasonable return to the public.



CASE STUDY 01

The University of Cambridge's Accelerate Science initiative



CASE STUDY 02

KU Leuven's Leuven.AI programme



CASE STUDY 03

The University of Oxford's AI in Science fellowships



CASE STUDY 04

Data Science Africa (DSA)



CASE STUDY 05

The Machine Learning for Science Cluster at the University of Tübingen



CASE STUDY 06

The Alchemy hub for AI in chemistry



CASE STUDY 07

The UK Centre for Ecology and Hydrology (UKCEH)



CASE STUDY 08

The European Centre for Medium-Range Weather Forecasts (ECMWF)



CASE STUDY 09

The OpenBind consortium



CASE STUDY 10

Polymathic AI

Building the missing middle in AI for science

A layer of institutional work sits between AI policy priorities or investments and their productive use in science. This layer can take different forms: for example, engineers who translate between AI and domain researchers; mechanisms for rapid prototyping and experimentation; community networks for peer learning; or processes for testing and verification. Together, they help make input-side infrastructure investments productive.

There is no single institutional model in the cases. Leuven built a small engineering service; DSA a distributed network; AIchemy a national hub. These reflect different local needs and constraints. Together they show certain core functions that allow the research system to absorb new AI capabilities (Table 2).

Supporting these capabilities also requires choices about how existing resources are allocated. Funders and institutions need to create space for emerging approaches, including by reviewing where established programmes, facilities, or technologies continue to deliver value. This requires mechanisms for learning which experiments should be extended, which capabilities should become sustained infrastructure, and which activities should be discontinued. Part of the policy challenge in AI for science is to develop a research funding system capable of reallocating resources as technologies change and new evidence about what works emerges.

Much of this work falls between conventional categories: it is not capital infrastructure, and it is not a traditional research project; it requires skilled people, but they are not necessarily conventional academic posts.

Table 2: Capabilities for AI absorption

Function	Question	How institutions provide it	Examples from cases
Explore	Where can researchers test emerging AI capabilities and discover useful applications?	Flexible and long-term funding; open calls and seed grants; access to experimental compute; protected time for interdisciplinary work	Tübingen, Polymathic, AIchemy
Translate	How is AI expertise connected to scientific problems?	Shared engineering services; embedded technical specialists; clinics and short exploratory projects; mixed AI–domain teams	Leuven, Accelerate, AIchemy, OpenBind
Diffuse	How does capability spread beyond researchers and institutions already at the frontier?	Responsive training; fellowships and cohorts; communities of practice; mentoring; distributed networks and reusable resources	Accelerate, Oxford, DSA, UKCEH, Tübingen
Equip	What makes AI practically usable for scientific work?	AI-ready data and data stewardship; appropriate compute; experimental facilities; models, software and interfaces that researchers can use and adapt	Polymathic, OpenBind, ECMWF, UKCEH, Accelerate, AIchemy
Assure	How can researchers establish when AI systems are good enough to rely on?	Experimental deployment; domain-specific testing and verification; blind challenges; reproducible workflows; scientific oversight and scrutiny	ECMWF, OpenBind, UKCEH, Tübingen
Embed	What allows a successful experiment to become a lasting part of scientific practice?	Integration into workflows; permanent technical roles; maintenance arrangements; institutional services; long-term funding for shared capabilities	ECMWF, UKCEH, Oxford, Tübingen, AIchemy

These functions point to a different model of how AI will change science. Investment in compute, models, AI talent, and new ventures aims to expand what AI can do. The translation of those advances into scientific progress depends on a process closer to technology diffusion: research institutions developing the skills, roles, services, and working practices needed to put new capabilities to use. In the process, operational experiences expose limitations that benchmarks miss and find directions to inspire a new wave of innovation in AI technologies.

These institutional functions also have implications for the design of research funding and policy. The case studies suggest a number of ways in which funders and policymakers could enable adoption across the research system (Table 3).

Table 3: Enabling AI adoption across the research system

AI for science challenge	Enablers and potential responses
Connecting researchers to AI capabilities	Build practical interfaces between national infrastructure and research communities, including supported onboarding, short engineering engagements, and shared technical services. Design these alongside infrastructure investment to cultivate demand.
Supporting experimentation and learning	Provide flexible funding for exploratory work and use adoption programmes to identify emerging scientific opportunities, unmet demand and recurring technical bottlenecks. Feed these insights into future research and infrastructure priorities.
Making choices about investment	Build mechanisms for reviewing investments as scientific needs and technologies change. Use evidence from experimentation and adoption to inform decisions about which activities to expand.
Building and retaining expertise	Develop funding and career structures that recognise engineering, data stewardship and interdisciplinary work. Support shared and distributed models that allow technical expertise to work across research groups and institutions.
Moving from pilots to sustained use	Create routes for the transition from experimental tools to dependable scientific services, covering domain-specific validation, user testing, integration into workflows, and long-term maintenance.
Preserving scientific autonomy	Explore long-term dependencies in funding and infrastructure decisions, including portability across compute environments, control over scientific data and models, and the ability to adapt or replace technologies. Consider how publicly funded scientific assets can continue to generate public value as AI capabilities develop.

Adoption is a process of experimentation and innovation. Universities, research institutions, national facilities, and research communities are addressing different versions of the adoption problem in different ways. This variation can be a source of policy evidence. Policy needs mechanisms to support experimentation and to learn from it.

The UK's scientific institutions are more than destinations for AI technologies developed elsewhere. They are places where the uses, limits, and value of AI technologies in science will be discovered, and future AI capabilities incubated. The first phase of AI for science policy has concentrated on building AI capability. The next needs to learn how that becomes scientific capability, considering what infrastructure will allow the research system to absorb new AI capabilities, while those capabilities continue to change.

CASE STUDIES

Thank you to case study participants for their contributions to this brief:

- Philipp Berens, Director of the Hertie Institute for AI in Brain Health, Professor of Data Science at the University of Tübingen
- Professor Gordon Blair, Head of Environmental Digital Strategy at UKCEH and Co-Director of the Centre of Excellence in Environmental Data Science (CEEDS)
- Matthew Chantry, Strategic Lead for Machine Learning, European Centre for Medium-Range Weather Forecasts
- Miles Cranmer, Assistant Professor in Data Intensive Science, Department of Applied Mathematics and Theoretical Physics and Institute of Astronomy, University of Cambridge
- Professor Ben Lambert, Director, Schmidt AI in Science Programmes, University of Oxford
- Dr. Godliver Owomugisha, Senior Lecturer & Lead, Busitema AI & Interdisciplinary Research Group, Busitema University
- Neil Lawrence, DeepMind Professor of Machine Learning, Chair, ai@cam University of Cambridge
- Ciira Maina, Director - Centre for Data Science and Artificial Intelligence (DSAIL), Dedan Kimathi University of Technology and Board Chair - Data Science Africa
- Jessica Montgomery, Director, Accelerate Programme for Scientific Discovery and ai@cam, University of Cambridge
- Frank von Delft, Professor of Structural Chemical Biology, University of Oxford and Principal Scientist at Diamond Light Source
- Sophia Yaliraki, Co-Investigator and Hub Ambassador, AIChemy and Professor of Theoretical Chemistry, Imperial College London

Thank you also to science writer Sarah Griffiths for her work with case study participants.

Footnotes

1. Zhao, C. (2026) 'OpenAI breakthrough triggers existential crisis in math', Science, 393(6817), 15 September. doi:10.1126/science.zbgmw3q.

2. Boboris, K. (2026) 'Fair moderation, equitable access, and AI: arXiv's updated rate limit policy', arXiv Blog, 1 October. Available at: https://blog.arxiv.org/2026/10/01/updated-rate-limit-policy/ (Accessed: 2 November 2026).

3. Bretter, C., Todorova, B., Curtis, C. and Hornsey, M.J. (2026) 'Public views on the use of AI in scientific research', Research Policy, 55(10), 105632. doi:10.1016/j.respol.2026.105632.

4. Codreanu, M. et al. (2026) AI in science: early insights. Available at: https://futuretech.mit.edu/publication/ai-in-science-early-insights (Accessed: 26 October 2026).

5. And projects such as Polymathic AI, featured in this brief.

6. Berens, P., Cranmer, K., Lawrence, N.D., von Luxburg, U. and Montgomery, J. (2023) 'Machine learning for science: bridging data-driven and mechanistic modelling (Dagstuhl Seminar 22382)', Dagstuhl Reports, 12(9), pp. 150-199. doi:10.4230/DagRep.12.9.150.

7. Cranmer, K., Lawrence, N.D., Montgomery, J. and Thérien, D. (2026) 'AI for science: reframing AI's role in discovery', RSS: Data Science and Artificial Intelligence, 2(1), udag003. doi:10.1093/rssdat/udag003.

8. Department for Science, Innovation and Technology (2025) AI for science strategy. Available at: https://www.gov.uk/government/publications/ai-for-science-strategy/ai-for-science-strategy (Accessed: 2 November 2026).

9. ai@cam (2025) Navigating AI sovereignty: strategic choices for the UK. Cambridge: ai@cam. Available at: https://www.ai.cam.ac.uk/reports/navigating-ai-sovereignty-strategic-choices-for-the-uk-2025/ (Accessed: 2 November 2026).



AI FOR SCIENCE:

Case Study 1

Accelerate Programme for Scientific Discovery, University of Cambridge

A community-led model for building AI capability in research.

KEY TAKEAWAYS

- **Human capital matters as much as hardware infrastructure.** Without people who can help researchers navigate the transition to GPU-based workflows, high-performance computing investment risks being underused or poorly matched to AI-enabled science. Policies that fund infrastructure without funding the accompanying human support, and institutions who fail to understand how to recruit and develop this human capital, will underdeliver.
- **Flexible, sustained funding enables an iterative, community-responsive model that rigid grant cycles cannot.** Accelerate has been able to pivot as the technology changes, for example launching a large language model workshop within weeks of demand emerging. Project-by-project funding structures make this kind of agility difficult to replicate.
- **Local conditions are key to embedding success.** Accelerate benefits from an ecosystem experienced in AI deployment. Policies that aim to scale this model should consider what these local conditions provide, and what substitute conditions would be needed elsewhere.

CONTEXT AND CHALLENGE

When the Accelerate Programme for Scientific Discovery was established at the University of Cambridge in 2020, AI was already beginning to transform research in narrow domains – protein structure prediction, image analysis, and natural language processing were all areas of early adoption. For most researchers, however, these advances were inaccessible. The skills required to work with machine learning were difficult to acquire; training provision was patchy and slow to update; and the cultural and organisational gap between AI and domain scientists remained largely unaddressed.

The result was: a small number of research groups – typically those already close to computer science or engineering – were pulling ahead, while the majority continued working without AI capabilities. The gap was in skills, in support structures, and in the research infrastructure that would allow AI expertise to diffuse across disciplines.

APPROACH AND IMPLEMENTATION

Training: responding to demand at speed

The core of Accelerate's model is a training offer designed to move faster than conventional curricula. When large language models became a priority for researchers, the team developed and launched a new workshop within weeks, alongside continuing provision on fundamentals such as Python, code packaging, and the deployment of diffusion models. A core team of four machine learning engineers has supported over 2,500 researchers across more than 90 departments, with more than 700 engaging with training programmes in 2025 alone.

This speed is enabled by two design choices. First, the programme does not attempt to anticipate all training needs from the centre: direction is shaped by demand from across the University, creating what the team describes as a 'supply chain of ideas' that matches researcher needs with available solutions. Second, the programme supplements its permanent team with external specialists – in areas such as industrial software engineering – when demand exceeds in-house capacity, allowing it to scale without the delays associated with permanent hiring.

The training offer is supported by an 'AI Clinic' that operates in two modes: AI Cafés providing informal drop-in spaces for quick consultations and cross-disciplinary exchange, and ongoing consultancy for more complex challenges. These convening mechanisms have evolved as the community has grown, and now function as a knowledge-exchange infrastructure in their own right.

Fellowships and seed grants

Alongside training, Accelerate hosted five-year fellowships for early-career researchers working at the intersection of AI and a scientific discipline. These provided sustained support for individuals developing interdisciplinary careers that are not traditionally well-supported in conventional academic incentive structures.

Small grants – ranging from £10,000 to £25,000 – have proved an effective mechanism for seeding community formation. Over 40 projects have been funded, spanning marine ecology, dementia research, and sustainable plastics. In 2025, 15 collaborative projects were supported, with several transitioning from proof-of-concept to substantive outcomes. The grants are notable not only for their research outputs but for the cross-disciplinary connections they create.

Bridging research to HPC infrastructure

The programme's relationship with Dawn – the national AI Research Resource (AIRR) supercomputer hosted in Cambridge – illustrates both the value of its bridging role and the policy challenges around infrastructure utilisation. Within four months of Dawn's launch, Accelerate had recruited 44 'pioneer projects' – research groups with AI for science needs – and given them access to the system. This provided immediate use cases for an infrastructure that might otherwise have sat underused while researchers found their way to it. It also gave Dawn's developers direct feedback on what scientists need from large-scale compute. Moving a research project onto GPU-based computing is not straightforward for researchers without prior experience of HPC environments; Accelerate's role in providing the human expertise to navigate these transitions is important in translating infrastructure into impact. Without Accelerate's connections to communities deploying AI in science and public services, University HPC infrastructure would have developed disconnected from the communities it should serve.

Building open AI infrastructure

Previous Cambridge innovations such as the Raspberry Pi and Acorn computers helped make computing accessible by both lowering cost and giving people the ability to experiment, learn, and build their own applications. Inspired by this history, Accelerate is exploring how a similar approach might support the adoption of AI in science: creating open, modular foundations that allow researchers to adapt and develop AI tools for their own scientific purposes, and avoiding dependencies on systems developed elsewhere.

This approach envisages an environment in which researchers can combine different models, scientific datasets, and computational tools to build AI-enabled workflows suited to their needs. Preconfigured components and support for modest computing infrastructure could lower barriers to entry, shared protocols and transparent workflows would support reproducibility, while a carefully-designed harnessing ensures safety of experiments. Instead of requiring researchers to adopt a single platform, the aim is to create common foundations that scientific communities can adapt and extend.

An early illustration is Denario, an open-source agentic system supported through Accelerate's seed funding, which brings together AI agents to support scientific tasks including hypothesis generation and research development (Borrett et al., 2026; Xu et al., 2025). While this is one component of a larger vision, Denario demonstrates how scientific workflows can be assembled from reusable components, providing a starting point for experimentation with more adaptable AI research tools.

The wider opportunity is for research institutions to develop shared AI capabilities that scientists can help shape from the ground up. Such an approach could enable innovation across disciplines while giving researchers greater control over their tools, data, and computational dependencies. It suggests a role for publicly supported research programmes in building access to AI through open technical foundations and communities through which new scientific applications can emerge.

Funding and institutional positioning

Accelerate is funded by Schmidt Sciences, a philanthropic organisation whose approach has been important to the programme's ability to operate as described. The flexibility to pivot, experiment and respond to emerging needs would be constrained under conventional grant structures, which typically

require defined outputs to be specified in advance.

This raises a structural question for policy. The kind of iterative, community-responsive model Accelerate has developed is difficult to specify in a funding application, because the value is in the ability to respond to conditions that are not yet known. Project-by-project funding cycles tend to reward programmes that can be pre-specified, which may systematically disadvantage models like this one.

IMPACT AND OUTCOMES

Accelerate has trained over 2,500 researchers across 86 departments, supported over 40 seed-funded projects, and created training materials reaching 4,500 users – substantial numbers for a team of four core staff.

The programme has also seen a shift in research culture: a growing normalisation of AI methods across disciplines that would not have happened through top-down mandates, enabled in part by wider curiosity about the potential of large language models and generative AI in science. The 'train-the-trainer' approach, through which researchers are equipped to share AI methods within their own communities, is designed to drive diffusion rather than dependence on Accelerate itself. Whether this can happen at scale – whether trained researchers are in practice sharing what they have learned in a self-sustaining way – is an important open question for the model's long-term value.

LESSONS FOR POLICY

Infrastructure investment requires investment in people

Compute infrastructure does not use itself. Accelerate's experience recruiting pioneer projects onto Dawn suggests that without concurrent investment in human capacity – people who can help researchers navigate GPU-based workflows and make productive use of HPC systems – the UK's AI Research Resource investment risks being underused, particularly by researchers in fields without prior HPC experience. While many HPC teams now identify the need for such interfaces, there is little evidence of such teams being developed around major UK infrastructure investments.

Flexible, sustained funding is structurally important, not incidental

Accelerate's effectiveness is not separable from its funding model. The programme's ability to respond rapidly to emerging needs – launching workshops in weeks, pivoting to new tools as the landscape shifts – depends on a funder willing to support experimentation without requiring pre-specified outputs.

The train-the-trainer model has real potential but unresolved limits

The ambition to equip every PhD student at Cambridge with AI skills, and to diffuse that capability organically through social networks, is plausible in the right conditions. Those conditions include a research community dense enough for informal knowledge exchange, a concentration of AI expertise to draw on, and a critical mass of trained researchers to sustain the network. The question for policy is whether comparable conditions can be created – or approximated – elsewhere.

The sovereignty question is live but unresolved

As agentic AI systems become standard research tools, decisions taken now about the infrastructure layer on which they run will determine who controls the data those systems learn from, who captures the value they generate, and whether the resulting capabilities remain accessible as public goods. Projects within the University such as AgentLab and Aardvark Weather (Allen et. al, 2025) show how the UK has both the technical capability and the institutional motivation to build sovereign research AI infrastructure. Whether it has the funding models and policy frameworks to support that work at scale is less clear.

References

Borrett, T., Xu, L., Nilipour, A., Bolliet, B., Pierre, S., Allys, E., Lecat, C., Dai, B., Chang, P.W. and Bhimji, W. (2026) 'Competing with AI scientists: agent-driven approach to astrophysics research', arXiv, arXiv:2604.09621. Available at: https://arxiv.org/abs/2604.09621 (Accessed: 2 November 2026).

Xu, L., Sarkar, M., Lonappan, A.I., Zubeldia, Í., Villanueva-Domingo, P., Casas, S., Fidler, C., Amancharla, C., Tiwari, U., Bayer, A. and Ekioui, C.A. (2025) 'Open source planning and control system with language agents for autonomous scientific discovery', arXiv, arXiv:2507.07257. Available at: https://arxiv.org/abs/2507.07257 (Accessed: 2 November 2026).

Allen, A., Markou, S., Tebbutt, W., Requeima, J., Bruinsma, W.P., Andersson, T.R., Herzog, M., Lane, N.D., Chantry, M., Hosking, J.S. and Turner, R.E. (2025) 'End-to-end data-driven weather prediction', Nature, 641, pp. 1172-1179. doi:10.1038/s41586-025-08897-0.



AI FOR SCIENCE:

Case Study 2

KU Leuven Institute for Artificial Intelligence (Leuven.AI)

An internal service providing focused, charged-for, time-capped projects that gets machine learning expertise to the domain scientists who need it.

KEY TAKEAWAYS

- **AI adoption often depends more on engineering capacity than on new AI research.** Many research groups do not need a permanent machine learning researcher; they need short bursts of expert engineering support to adapt existing methods to their scientific problems. Leuven shows that research software engineers can provide this capability at institutional scale. Funders and research institutions should recognise and incentivise research software engineering as a distinct career track.
- **Small, high-volume AI support can have wider impact than high-profile projects alone.** By limiting projects to €10,000–€30,000 and no more than three person-months, Leuven spreads AI capability across many departments that would never access large AI funding programmes. AI for science funding should therefore include mechanisms that support many small interventions alongside major strategic investments, and grant rules should treat research software engineering provision as institutional capacity rather than external consultancy.
- **Shared AI services become strategic institutional infrastructure.** Beyond delivering individual projects, Leuven's AI as a Service (AIaaS) gives the University a picture of where AI demand is emerging, what infrastructure researchers need, and where investment will have the greatest impact. Shared engineering services create institutional capability as well as project outputs.

CONTEXT AND CHALLENGE

AI adoption in science is often framed as a shortage of algorithms, data, or compute. Leuven.AI began from a different diagnosis: across much of the University, the missing ingredient was access to experienced machine learning engineers. KU Leuven established Leuven.AI in 2020 to coordinate AI activity across a university that already had internationally recognised AI researchers, but whose expertise was siloed inside individual labs. Domain scientists in biology, materials, the social sciences and elsewhere who wanted to use machine learning faced recurring problems around access to expertise and resources. Without experience or training, they struggle to deploy machine learning themselves and face challenges in convincing funders of their work's viability. Collaborations with AI researchers are often not feasible: applying established AI methods to another discipline's problems does little to advance an AI researcher's own publication record and career, so the incentives to take on such collaborations can be limited. The academic hiring cycle can be too slow for short engagements.

Implementation and adaptation of existing AI methods – together with good software practice and effective user interfaces – remains a significant challenge for many research teams. Leuven.AI's response was to build an internal service, branded 'AI as a Service' (AIaaS), that domain researchers can commission for short, paid engagements. The name borrows from cloud-industry usage but the substance is different: this is internal research software engineering, billed to research groups and delivered by a dedicated team that sits inside the University. The service is deliberately small, and built for reach.

APPROACH AND IMPLEMENTATION

A cost-recovery model that doubles as a demand filter

AIaaS launched on one year of seed funding from the University to hire a single machine learning research engineer. From there, it was designed to be self-sustaining. Research groups are charged at 1.5 times personnel costs: enough to cover the engineer's salary plus a 50% margin that funds non-billable management time and gradually accumulates capital for additional hires. The pricing is the central design choice. Charging filters out tentative or low-commitment requests – a principal investigator who allocates real budget treats the project as a real project, whereas free central support tends to attract requests that are abandoned partway through. The same charge that filters demand also funds growth.

The model has limits. The 1.5x markup is incompatible with EU framework programme grants such as Horizon Europe, which permit recovery of personnel costs but not the additional margin. The cost-recovery formula makes the service self-funding, but it also limits which grants can pay for it.

Optimising for breadth

AIaaS projects are capped at three person-months, enabling one AIaaS expert to work on approximately ten projects per year. A common entry point is a one-month 'exploratory engineering' engagement designed to test whether an idea is viable before larger commitments. The cap is a design feature: it lets a small team stay responsive to a steady flow of new requests rather than disappearing into long-running collaborations. It also creates a boundary with academic AI research. Problems requiring deeper algorithmic novelty are handed back to KU Leuven's AI researchers, while AIaaS handles the engineering work that does not need new methods but does need someone with high-quality AI and software engineering skills to ship a tool or service.

Typical project budgets sit at €10,000–€30,000. AIaaS is not trying to enable foundation-model-scale work; it is trying to make small machine learning projects feasible at high volume, reflecting a judgement about where the gap in the local research ecosystem lies.

Building capability, not dependency

Leuven.AI encourages interdisciplinary cooperation through a range of engagement routes, from informal roundtables and an 'AI Café' for 30–50 people on specific topics to large conferences attracting over 400 participants. These events are essential for sharing knowledge across KU Leuven's diverse faculties, but there are limits to how many enduring collaborations form directly from talks and conferences alone. The AIaaS engineers are in place so that domain scientists can incorporate machine learning into their workflows at the moment they need it.

AIaaS offers three engagement types: consultations (advice on algorithm choice without code delivery), co-development (the AIaaS engineer builds tools for the domain scientist's problem), and brokered collaborations with KU Leuven AI researchers when off-the-shelf tools are insufficient. All three are structured around daily-to-weekly contact between the AIaaS engineer and a named person in the requesting lab – typically a PhD student or postdoc – to ensure knowledge transfers rather than being delivered as a black box. Every project closes with a written report and code published on an internal GitLab accessible to colleagues. AIaaS claims no ownership or equity in any output: results remain with the commissioning research group, which means researchers are willing to use the service for work they intend to commercialise.

Challenging assumptions about compute bottlenecks

AIaaS's experience challenges two common assumptions about compute. The first is that every AI for science project needs high-performance computing. In practice, many projects – particularly those not involving computer vision or large language models – run comfortably on standard university computing, with the Flemish supercomputer reserved for the minority of cases that need it. The second assumption is that generative AI and large language models require large compute, typically provided via third-party cloud services. Together with KU Leuven's ICT services (ICTS), Leuven.AI set up open-weight locally hosted large language models, removing infrastructure and implementation burdens from individual labs, while speeding up application development by allowing researchers to focus on domain-specific parts of models. The set up allows researchers to use large language models on proprietary and sensitive data, ensuring that it stays inside the institution, reducing risk associated with using external tools.

The team has found that combining domain knowledge with task-specific tuning achieves useful performance with relatively small open-source models, and is monitoring what new infrastructure the continued rise of generative AI will require.

IMPACT AND OUTCOMES

AIaaS delivers benefits at three levels. At the level of individual projects, it enables research groups to incorporate machine learning methods that would otherwise be beyond their capacity, delivering code, prototypes and technical expertise through short, focused engagements. At the level of research groups, it strengthens grant applications by helping researchers incorporate credible AI components into funding proposals

without recruiting permanent machine learning specialists.

At the institutional level, AIaaS provides something few universities possess: a university-wide view of where demand for AI is emerging. Because researchers from across the institution approach the service with their technical challenges, it identifies recurring bottlenecks, infrastructure needs and skills gaps across disciplines. This evidence informs institutional planning, including decisions about central AI infrastructure and researcher support.

Demand is already growing beyond the capacity of a service, which was initially operating primarily with a single engineer. This is evidence of its value, but also of the challenge of scaling it. Its required margin limits growth, being too low to scale organically or too high for some major funding schemes, limiting access for researchers whose grants cannot recover the required margin.

The three-month cap keeps the service responsive to many groups. While short-term projects can be a limitation, a large base of academic AI researchers enables projects requiring sustained methodological development to transfer from AIaaS to labs for longer-term collaborations. The model enables valuable flexibility for researchers, but it comes at a financial cost, as the administrative effort required to set up and manage projects accounts for a share of the margin.

While the operating model and financial structure of the service could be replicated across other universities, key to the success of Leuven.AI is the way it is embedded across the institution's ecosystem and strong machine learning research base.

LESSONS FOR POLICY

Cost recovery can sustain small services, if funding rules permit it

Leuven.AI's 1.5x cost-recovery formula is what makes AIaaS self-sustaining after seed funding runs out. The same formula is incompatible with Horizon Europe and similar schemes that cap recoverable costs at personnel level. The current EU situation shows what happens when funding rules and service models are out of sync: the service exists and is taken up, but a meaningful fraction of researchers cannot access it through their main grants.

Fund research software engineering as a tier above project hiring

AIaaS exists because the academic hiring cycle and the publication-driven incentive structure cannot deliver short, high-quality engineering on demand. Research software engineers sit between these two systems and need their own funding tier. A workable model is layered: large, machine-learning-intensive labs employ permanent in-house engineers, while a centralised flexible service handles smaller or less technical departments. Institutions and funders should expect both tiers are needed.

Pricing, time caps and no-equity rules underpin service design

What makes AIaaS work is contained in three design choices: the cost-recovery formula that filters demand and funds growth; the three-month cap that keeps the team responsive and protects the boundary with academic research; and the no-equity policy that keeps results with the commissioning group, lowering barriers to commercialisation and spin-out routes handled elsewhere in the University. Each of these design choices is load-bearing. Removing any one of them changes what the service is.

Build administrative AI literacy alongside research AI capacity

Research support offices that handle complex grant applications now need their own AI literacy: to advise researchers on the credible use of generative tools in proposals, to assess AI components in submitted bids, and to engage with funder requirements that increasingly reference AI methodology. An institution with strong lab-level AI capability and weak administrative AI literacy will under-deliver on its grant pipeline.

AI FOR SCIENCE:

Case Study 3

AI in Science postdoctoral and faculty fellowships at the University of Oxford

Investing in researchers, not just projects: fellowship cohorts as AI capability infrastructure.

KEY TAKEAWAYS

- **Building AI capability requires investment in people.** Oxford's fellowship programmes combine flexible funding, mentoring, research software engineering support, and peer networks to develop researchers who can apply AI confidently throughout their careers.
- **Research capability develops through communities as well as formal training.** Carefully designed cohorts create durable peer networks that accelerate knowledge sharing across disciplines, support career development, and continue generating collaborations long after the fellowship ends.
- **Engineering support should be embedded within researcher development.** Research software engineers do more than provide technical assistance. By teaching robust software engineering practices alongside AI methods, they help researchers produce more reproducible science while developing skills that outlast individual projects.

CONTEXT AND CHALLENGE

Governments seeking to accelerate AI in science often focus on compute, data infrastructure, and research funding. These investments are essential, but not sufficient. Many researchers lack the opportunity to develop the technical skills, software engineering practices and interdisciplinary networks needed to incorporate AI confidently into their work. Building AI capability therefore requires more than funding projects: it requires sustained investment in people.

These challenges are sharpest where resources are scarcest. Access to the infrastructure, skills and collaborative networks needed for AI-enabled science is distributed highly unevenly across the world, yet the expertise and context-specific knowledge of scientists working on global challenges such as climate change is essential to keeping AI models useful and grounded in reality. Building equitable, collaborative ecosystems in which AI expertise can diffuse across disciplines and borders benefits both sides of that exchange.

The University of Oxford's AI in Science fellowship programmes were designed around this challenge. Rather than funding narrow research projects, they invest in researchers as scientific practitioners. Fellows receive flexible funding, individual mentoring, research software engineering support, advanced AI training, and membership of carefully designed interdisciplinary cohorts. The programmes operate at two career stages: the Eric and Wendy Schmidt AI in Science Postdoctoral Fellowship supports researchers across Oxford applying AI within STEM disciplines, while the Schmidt AI in Science Faculty Fellowship supports early-career faculty from Africa and India, combining AI capability building with international research partnerships.

APPROACH AND IMPLEMENTATION

Funding model

The fellowship programmes prioritise researcher development over tightly specified project delivery. Rather than requiring predefined milestones or measuring success primarily through publication counts, fellows have the flexibility to pursue the outputs that best advance their research, including papers, software, datasets, or practical tools.

This flexibility is enabled by philanthropic funding from Schmidt Sciences. The Postdoctoral Fellowship, launched in 2022 as part of a £148 million global programme across nine universities funding approximately 160 postdoctoral fellows and 20 faculty fellows internationally each year, provides at Oxford a six-year commitment supporting around 55 postdoctoral fellows, covering salary together with research and travel funds. The Faculty Fellowship Programme, launched in 2025, supports 18 early-career faculty members over five years, focusing on researchers from Africa and India working on climate change and biodiversity, food and water security, and STEM underpinning biomedical research. The deliberate absence of tightly specified milestones creates space for activities that are difficult to justify through traditional funding mechanisms, particularly mentoring, interdisciplinary cohort building and sustained researcher development.

Governance

Both programmes operate within Oxford's formal governance structures but are managed with local autonomy. Day-to-day delivery is led by a dedicated Oxford team supported by an Academic Steering Group, while Schmidt Sciences maintains light-touch oversight focused on strategic progress. This allows the programme team to adapt training, mentoring and cohort activities as fellows' needs evolve, rather than delivering against rigid externally defined milestones.

For the Faculty Fellowship Programme, governance extends through partnerships with the Africa Oxford Initiative (AfOx) and the Oxford India Centre for Sustainable Development (OICSD). These partnerships bring regional expertise directly into recruitment, mentoring and programme design, helping ensure that capability building reflects the institutional contexts in which fellows will continue their careers, rather than simply reproducing an Oxford model overseas.

Cohorts as communities of practice

The programmes assume that AI capability develops through sustained interaction between researchers, not through formal training alone. Rather than treating fellows as independent award holders, both programmes deliberately create interdisciplinary cohorts that function as long-term communities of practice, enabling researchers to exchange AI methods across scientific domains, develop collaborations, and support one another's professional development beyond the fellowship itself.

For the Postdoctoral Fellowship, this happens through regular research colloquia, professional development workshops, and peer-led activities that create repeated opportunities for interaction. Fellows take a peer approach to career development: sharing grant and job applications for feedback, providing exemplars of past funded projects, and running mock interviews within the cohort in a supportive environment. The Faculty Fellowship Programme builds cohorts differently. Fellows are selected within three thematic areas, creating immediate opportunities for collaboration around shared scientific challenges, and fostering enduring research networks between participating institutions across Africa, India, and the UK.

Mentoring and research software engineering

The programmes treat AI capability as something learned through practice. Each fellow receives a combination of technical mentoring, individual AI training, and hands-on software engineering support tailored to their research. The design has already adapted to experience: the Faculty Fellowship Programme initially planned shared lectures and workshops covering common AI topics, but early experience showed fellows' backgrounds and research needs were too diverse for a standard curriculum. The programme shifted to individually tailored training, combining graduate-level courses with one-to-one tuition that develops AI knowledge alongside each fellow's research programme.

The Postdoctoral Fellowship combines disciplinary and technical mentoring. Fellows are embedded within a principal investigator's research group while also being paired with a senior AI faculty mentor, ensuring that scientific questions and AI methodology develop together rather than separately. Research software engineers complete the model. They teach fellows how to implement AI methods using robust software engineering practices that produce reusable, reproducible research through intensive project-based training on arrival, followed by continuing access to engineering expertise as each fellow's research evolves.

Infrastructure as enabler

The programmes provide access to Oxford's wider AI and research computing infrastructure, but treat infrastructure as an enabler rather than the primary intervention. Fellows are encouraged to use centrally supported university AI

services wherever these meet their needs, with commercial AI tools made available where justified by individual research projects, alongside guidance on their appropriate use. The aim is to balance researcher productivity with responsible AI use, institutional oversight and long-term sustainability.

IMPACT AND OUTCOMES

The programmes aim to create three complementary forms of impact. First, they enable research that would otherwise be difficult to undertake, by equipping scientists with AI methods, engineering support and the freedom to pursue ambitious interdisciplinary ideas. Second, they build lasting capability: fellows equipped with research software engineering skills can create software artefacts that act as impact multipliers throughout their careers. Third, they establish enduring peer networks that continue generating collaborations and knowledge exchange beyond the lifetime of individual fellowships.

These forms of impact are becoming visible. Fellows have initiated cross-disciplinary collaborations, including a publication exploring transparent and reproducible systematic literature reviews using AI methods, delivering work unlikely to have emerged through conventional disciplinary funding. Fellows are also increasingly transitioning into academic leadership roles and equivalent roles outside academia, supported by peer mentoring, career development activities and continuing engagement with subsequent cohorts; transitional funding enables alumni to remain active within the community, strengthening cumulative impact over time.

LESSONS FOR POLICY

Flexible funding enables capability building

Oxford's fellowship programmes use flexible funding to give fellows time to acquire new technical skills, work across disciplinary boundaries, receive mentoring and engineering support, and pursue outputs beyond conventional publications. These investments are difficult to accommodate within tightly specified grant schemes but are central to building lasting AI capability. Public funding programmes should consider whether greater flexibility is needed to support researcher development alongside project delivery.

Treat researcher communities as research infrastructure

Structured cohorts, repeated opportunities for interaction, peer mentoring, and continued alumni engagement create communities that exchange methods across disciplines, accelerate career development and generate collaborations that conventional project funding rarely produces. These communities become a form of research infrastructure, enabling capability to diffuse through the wider research system long after individual fellowships end. Funding mechanisms should recognise and resource cohort building as a core component of AI capability, not an optional networking activity.

International capability building depends on institutional partnerships

The Faculty Fellowship Programme is designed to be mutually beneficial and to strengthen research capability within fellows' home institutions. Its 1+1+1 structure preserves fellows' permanent academic positions, funds teaching buyout before and after their stay in Oxford, and builds collaborations between institutions rather than dependence on a single university. International AI for science programmes require mechanisms that sustain capability locally, rather than focusing solely on researcher mobility.

Embed engineering expertise within researcher development

By teaching robust software engineering practices alongside AI methods, research software engineers enable researchers to produce reproducible, reusable computational research while developing skills that remain valuable throughout their careers. Embedding engineering expertise directly within researcher development programmes creates both immediate research benefits and long-term capability.

Long-term capability building currently leans on philanthropy

The flexibility that defines these programmes is possible because a philanthropic funder is willing to support researcher development without pre-specified outputs. That raises a question for public funders: if sustained, people-centred capability building is left largely to philanthropy, national AI capability priorities will in part be set by philanthropic rather than public choices. Public funding schemes that can accommodate this kind of investment would reduce that dependence.



AI FOR SCIENCE:

Case Study 4

Data Science Africa

A grassroots organisation designed to leverage data science and machine learning for the development of solutions to African problems

KEY TAKEAWAYS

- **Community-led networks can extend a small central team's reach.** DSA's summer schools drew over 1,800 applicants for 400 places in 2026, and its network is sustained largely by volunteers at partner universities. The model shows what a distributed, community-run structure can achieve on limited resources. It also highlights the scale of demand at grassroots level in research for engagement with AI.
- **Decentralised, community-led governance enables faster adaptation.** DSA's local chapters at partner universities set their own training priorities, often guided by participant surveys, and can respond to shifts in the field — such as new techniques, tools, or problems — without waiting for central approval. This is a structural advantage over curriculum-led capacity building. The trade-off is that the model depends on volunteer capacity and host-institution goodwill rather than institutionally-sustained funding.
- **Partnerships bridge compute gaps, but sovereign infrastructure is needed.** External collaborations provide vital access to compute, fellowships, and funding. However, national and regional investment in computing infrastructure is essential to unlock large-scale model development and fuel the ambitions of students and researchers working on public interest AI, especially in areas without immediate commercial return that might otherwise be neglected.
- **Sovereignty in AI for science has different dimensions.** Building regional AI capacity requires sovereignty over compute infrastructure (whether researchers can run the models their problems demand), sovereignty over the research agenda (whether they can direct what gets built), and sovereignty over where capacity accumulates (whether the researchers a programme develops build institutions in the region or are recruited out of it).

CONTEXT AND CHALLENGE

Well-funded international programmes play a vital role in supporting AI research across Africa, but they fall short of meeting the considerable demand from universities and their ambitious students. Many African researchers face a structural bottleneck: limited exposure to cutting-edge machine learning techniques, an absence of platforms designed to amplify research focused on African problems, and severe compute constraints. Most African universities can only afford lower-tier GPUs, which while sufficient for small-scale experimentation, lack the processing power needed to train large-scale models tailored to regional challenges such as smallholder farming. Consequently, building AI capability requires sustained investment and access to AI-related opportunities at a national and local level.

Data Science Africa (DSA) exists to create space for African researchers to lead work that is both locally relevant and globally significant. Founded in 2015, the grassroots organisation provides research funding, plus training in data science, machine learning, and emerging technologies, with the goal of leveraging these disciplines to develop solutions to Africa's unique challenges. Primarily focused on students at both undergraduate and postgraduate levels, it aims to prioritise capacity building and learning opportunities for African researchers and practitioners while emphasising practical applications of data science that can make a positive impact by providing AI-led solutions to local issues.

APPROACH AND IMPLEMENTATION

Funding and governance

DSA is funded through a hybrid model that combines tech industry partnerships, international development research grants, and academic host support. The host-institution model is crucial for DSA's success, with participating universities co-funding and hosting annual DSA summer schools as well as providing physical infrastructure, local logistics, and administrative support. Its model and blend of funding allow DSA to remain agile and address needs with specific programmes and research funding.

Arm has supported DSA's programmes, including sponsoring the summer school, providing expert input to teaching and programme development, and donating hardware necessary to train participants. DSA also worked with Google DeepMind in delivering its AI Research Foundations curriculum to help strengthen AI capacity across the continent. A total of 15 partner universities across Uganda, Kenya, Tanzania, Rwanda, Zambia, Cameroon, Botswana, Ghana, Zimbabwe, Malawi, and Nigeria oversaw implementation and mentorship, with Google DeepMind funding travel for selected students to attend DSA's flagship summer school hosted at Makerere University in Uganda in 2026.

Delivering targeted research

DSA's board, comprised of five directors, six members, and two advisors, ensures that funding is directed to the most critical areas such as healthcare and agriculture, with specific fellowships such as the DSA Women's Health (DSAWH) fellowship. Funded by Boston University through a grant from Bill and Melinda Gates Foundation, it supports 44 fellows and is designed to strengthen women-led research, expand data science capacity, and advance evidence-based solutions to critical women's health challenges across the continent.

The Llama Impact Grant for Startups and Researchers supports entrepreneurship, by giving five projects $20,000 in funding, technical mentorship, and access to a network of policymakers, and ecosystem partners. The 2025 cohort's projects spanned healthcare to agriculture, education to digital accessibility and demonstrated the transformative potential of open-source AI tools.

Flagship events for knowledge dissemination

DSA's primary activities include organising summer schools and workshops that equip participants with theoretical and practical skills in data science and provide an environment for network building and future collaboration.

The first event in Nyeri, Kenya in 2015 established DSA's event format of combining training with research presentations and has since been replicated in multiple countries across Africa, including Uganda, Tanzania, Nigeria, Ethiopia, Ghana, South Africa, and Rwanda. Today, demand far exceeds capacity. In 2026, more than 1,800 applicants spread across 65 countries competed for 400 places. The event remains financially lean, with accommodation provided by the university and processes overseen by volunteers at participating universities giving them experience in project management and leadership.

Building open AI infrastructure

Since its inception, DSA has evolved into a community that bridges the gap between academia, industry, and policymakers. This ecosystem is crucial, because while many students take computer-related courses, they do not always have knowledge of AI. To address this, local chapters largely run by volunteers at partner universities run small events to build capacity and motivate other universities to join DSA's network. To ensure topics and themes of events and training sessions are relevant, DSA is guided by the outcome of surveys in which future participants rank options.

Strategic collaborations with universities and local organisations provide the compute and other resources training sessions require. These partnerships help ensure that AI and data-driven solutions are grounded in local needs and contribute to addressing real-world challenges across Africa. DSA has built working relationships with government agencies, research institutes, and NGOs whose problems its participants work on. Engagement with private companies has been more challenging: firms have often not seen commercial value in AI applications aimed at local problems – smallholder agriculture, maternal health, or low resource languages, for example – where affected populations are not likely to generate revenue.

Another way in which DSA is deepening its regional engagement is through the formalisation of its Uganda chapter, which is a flagship component of its decentralised, community-driven model. The Uganda chapter offers regular meetups and convenings on topical issues, targeted training programmes for students and early career professionals, hackathons and training initiatives as well as grant writing and research dissemination workshops. It prioritises engagement through collaborative project design and implementation, and plans to produce an annual 'State of Data Science in Uganda' report. By embedding these activities within a locally led structure, the Uganda Chapter provides a scalable model for future DSA country chapters.

IMPACT AND OUTCOMES

Since DSA began a decade ago, it has trained hundreds of researchers, convened participants from dozens of African countries, and supported a growing portfolio of fellowships,

research grants and collaborative projects. In 2025, 351 participants attended its flagship summer school, while its hackathon demonstrated appetite for African-led innovation, attracting 1,757 participants. Winning projects included AI-assisted breast cancer screening, a Swahili speech recognition tool, and civic legal chatbots.

The DSAWH brought together 44 early-career women researchers from west and central Africa in two cohorts: 20 in 2025 and 24 in 2026. Angela Esi Ackon, a 2025 Fellow, published a study examining the economic burden of maternal mortality on households, communities, and national economies, providing policymakers, health systems and development partners with evidence that can inform effective investments in maternal health and contribute to the UN's third Sustainable Development Goal. A new DSAIL and Sunbird AI Fellowship - focused on advancing edge AI for biodiversity monitoring through bird sound classification in natural habitats - resulted in an optimised bio-acoustic sensing pipeline for Raspberry Pi devices, and a model able to classify bird species in the Dedan Kimathi Wildlife Conservancy with 98% accuracy.

As well as tools born from hackathons and grants, the self-sustaining community has created collaborations and opportunities for participants, with DSA alumni leading research at global tech companies including Google and Microsoft and African universities, as well as founding startups and shaping public policy.

LESSONS FOR POLICY

Building national AI capability requires an educational pipeline

DSA's summer schools, fellowships and hackathons reach hundreds of researchers a year. The oversubscription to summer schools indicates that demand for AI training substantially exceeds what current grassroots provision can meet. Sustained national capability requires expanded postgraduate provision — such as dedicated master's and PhD pathways — and earlier-stage AI content in secondary and undergraduate curricula. Centres of excellence need long-term funding to convert individual capacity into institutional capacity, but only where funding is long-term enough to offer the salaries, research infrastructure and career pathways that compete with international recruitment. Short-cycle or project-based funding constrains their growth and continuity.

Sovereignty in AI for science requires investment across multiple layers

National or regional compute infrastructure, accessible to higher education networks, would let African research institutions work at a scale that individual universities cannot support on their own. Alongside this, securing regional capability requires investment in: the skills to use AI, the ability to set research directions independently, and funding routes for the kinds of research the market does not reach. Without those, compute investment produces skilled participants in AI systems designed elsewhere.

Public–private alignment is essential to turn research into local economic value

Grassroots educational initiatives are effective at building academic talent, but their translation into local economic and social impact depends on whether the applications they enable get taken up and sustained. DSA's experience is that private-sector engagement has often been harder to secure for the applications with the clearest social value due to a lack of commercial returns. Policy frameworks that offset this market failure, through procurement rules, matched funding for public-interest applications, or dedicated public-interest AI funds, would let grassroots-developed talent build tools with the people they were designed for.



AI FOR SCIENCE:

Case Study 5

The Cluster of Excellence 'Machine Learning for Science' at the University of Tübingen

What 14 years of flexible national funding makes possible at the machine learning–science interface.

KEY TAKEAWAYS

- **Long-term, flexible funding enables different research than project grants.** The Cluster's 7+7 year commitment with lightweight oversight allows researchers to adapt priorities, explore unexpected opportunities, and build sustained interdisciplinary collaborations.
- **Embedded ethics expertise can shape research from inception rather than audit it retrospectively.** The Cluster employs philosophers, anthropologists, and legal scholars who work alongside machine learning researchers from project start, enabling ethical considerations to influence research design.
- **International partnerships built into large initiatives can serve strategic purposes beyond research collaboration.** The AIMS–Tübingen fellowship brings African researchers to Germany for training, and is now establishing a research presence in Rwanda through a joint chair – a move toward building sustained capacity rather than one-way knowledge transfer.

CONTEXT AND CHALLENGE

The Excellence Strategy, launched by German federal and state governments in 2006 and reshaped in 2019, aims to strengthen the international competitiveness of German science through sustained investment in research excellence. It funds Clusters of Excellence – collaborative research projects in internationally competitive fields – for initial seven-year periods, renewable for another seven. Applications require roughly 100 pages and a year of preparation, and around a quarter are funded.

The University of Tübingen had existing strength at the machine learning–science interface, with groups in neuroscience and the life sciences already using AI methods, plus a concentration of expertise through the Max Planck Institute for Intelligent Systems and the emerging Cyber Valley initiative. A group led by Philipp Berens and Ulrike von Luxburg proposed extending this approach across further disciplines, creating a Cluster of Excellence on Machine Learning for Science. The proposal succeeded, securing funding from 2019 with renewal through 2032.

APPROACH AND IMPLEMENTATION

Funding model

The German government's overall investment is substantial: 70 Clusters of Excellence share €539 million in 2026. Tübingen is one of eleven Universities of Excellence; it received initial funding for three Clusters in 2019, and three more were awarded funding in 2026.

Clusters receive funding for two seven-year periods. The model differs from shorter-term Collaborative Research Centres, which typically fund PhD student and postdoc salaries for up to four years with up to two renewals. The primary benefit is stability – funding secured for significantly longer – but Cluster funding is also more flexible. Control mechanisms are comparatively lightweight, enabling adaptation to new developments over each seven-year cycle and the exploration of unexpected interdisciplinary opportunities.

The Machine Learning: New Perspectives in Science Cluster, awarded in 2019, has over 70 members and their research groups across disciplines including neuroscience, linguistics, geosciences, educational sciences, biology, cognitive sciences, medicine, physics, environmental science, and economics. Partner institutions include the Max Planck Institute for Intelligent Systems and the Leibniz Institut für Wissensmedien, both in Tübingen, and the Cluster is part of the Cyber Valley ecosystem.

Enabling interdisciplinary research

The Cluster combines AI and machine learning with disciplines such as neuroscience, linguistics, physics and the geosciences, complemented by systematic ethical reflection on machine learning's impact on science and society. Four research areas were identified where progress is urgently needed: developing machine learning algorithms that reveal scientific laws from data; validating complex models and quantifying uncertainty; providing methods that let scientists control all stages of the machine learning workflow; and investigating how these developments affect scientific practice and the evaluation of evidence.

To drive progress in these areas, the Cluster funds professorships and independent research groups for early-career researchers, providing positions with greater stability and autonomy than standard postdoctoral roles. Notably, it offers fall-back permanent positions for early-career group leaders upon positive evaluation if they do not obtain a professorship elsewhere – a model that is unique in Germany, and a direct answer to the precarity that pushes AI talent out of academia.

Embedding ethics in research practice

The Ethics & Philosophy Lab, established within the Cluster during the first funding cycle, conducts ethical and philosophical research inside the scientific community rather than in a siloed discipline of its own. Now in its second phase and renamed the Norms & Practices Lab, it brings together researchers from philosophy, digital anthropology, science and technology studies, and law: a legal scholar, a cultural anthropologist and a philosopher funded by the Cluster work alongside machine learning PhD students, investigating how machine learning technologies change the norms and practices of science.

For the second funding phase, responsible AI has become a core priority. Where ethical questions were previously addressed ad hoc, the Cluster is formalising the approach through a dedicated Ethics Board, enabling responsible-AI considerations to be embedded at project inception. Demand is evident at every level, from PhD students to senior faculty; a central source of ethics expertise lets researchers navigate the landscape with professional guidance.

Supporting collaboration and skills transfer

The Machine Learning ⇄ Science Colaboratory brought together researchers and engineers to train, advise and collaborate with domain scientists at the Tübingen AI Centre. This initial approach to supporting AI deployment struggled to get off the ground in the context of the COVID-19 pandemic, which forced interactions online and reduced appetite for workshops. The experience highlighted the importance of understanding the different needs of user communities with different levels of AI literacy. Building on it, the Cluster is establishing an AI & Software Hub providing a physical space where AI experts offer implementation guidance.

The University offers machine-learning-affiliated programmes alongside courses in business, economics and neuroscience; the International Max Planck Research School for Intelligent Systems runs annual boot camps. An annual faculty retreat addresses machine learning requirements across research areas such as quantum physics, with a flat structure encouraging faculty to organise and lead workshops rather than attend top-down conferences.

Supporting internationalisation

The AIMS–Tübingen Data Science & AI Fellowship Programme, established in 2022 for a four-year period, offers five graduates of the African Institute for Mathematical Sciences (AIMS) annually a fully funded nine-month research fellowship in Tübingen. Fellows work on projects co-supervised by Cluster faculty and AIMS researchers, on problems ranging from climate network analysis (developing uncertainty estimates for climate network construction) to gravitational wave inference (evaluating machine learning methods for bias) and medical data science (exploring contrastive learning for medical image embedding).

The partnership has evolved beyond fellowships: the second funding period (2026–2032) includes a joint research chair in Machine Learning for Science co-hosted by the AIMS Research and Innovation Centre in Rwanda. This is a move toward building sustained research capacity in Africa rather than one-way knowledge transfer, though long-term outcomes depend

on where fellows establish their careers. The partnership serves both sides: advanced training and research connections for AIMS graduates in a field where local expertise is scarce, and where urgent needs like climate change adaptation and health system challenges could be amenable to AI, and extended international research networks for Germany.

Compute infrastructure and data governance

Cluster funding terms make it difficult for universities to pay for commercial cloud computing, so the University operates its own GPU cluster through its Training Centre for Machine Learning, alongside platforms such as EuroHPC. While funding for hardware procurement is ample, recruiting and retaining the expertise to manage it remains challenging, as specialists command higher salaries in the private sector, though recent shifts in the tech job market may ease the gap.

Cluster statutes require that data and models produced by members be released as open source, aligning with scientific norms of accessibility and reproducibility. German law, however, stipulates that inventions made by employees of public institutions belong partially to the institution, creating a potential tension. This grey area has yet to be tested; Cluster members are encouraged to embrace the open-source model to ensure broad scientific accessibility.

IMPACT AND OUTCOMES

The clearest evidence of impact is structural. From sports science to computational linguistics, Tübingen departments are independently investing in faculty with AI links. This suggests the Cluster has shifted the University's long-term trajectory, establishing AI as integral to scientific inquiry rather than a peripheral tool. The renewal of the Cluster through 2032, in a competition where only around a quarter of applications succeed, is itself an external validation of the first funding period.

The Cluster's design experiments are also generating transferable lessons: the fall-back permanent positions model for early-career group leaders, the pivot from a standard support unit (the Colaboratory) to a physical AI & Software Hub in response to what did and did not work, and the evolution of the AIMS partnership from fellowships to a jointly hosted chair. The University's wider success in Excellence competitions – six Clusters, the most of any German university – provides favourable institutional context, though it is an indicator of the University's standing rather than an outcome of this Cluster specifically.

LESSONS FOR POLICY

Long-term, flexible funding enables different research than project grants

The Cluster model provides 7+7 year funding with lightweight oversight, in contrast to typical three-year project grants requiring detailed plans and regular accountability. This enables researchers to adapt priorities as fields evolve, explore unexpected opportunities, and build sustained interdisciplinary collaborations without justifying every pivot.

Balance application rigour with operational flexibility

The Cluster application process is demanding – roughly 100 pages and a year of preparation. But once awarded, funding comes with lightweight oversight, giving researchers freedom to adapt over seven-year periods. The lesson is not simply 'reduce bureaucracy': rigorous selection can be justified if it secures long-term, flexible funding.

Embedded ethics expertise can shape research from inception

The Cluster's evolution from ad hoc ethical reflection to a dedicated Norms & Practices Lab and Ethics Board demonstrates one model for integrating responsible AI. Rather than treating ethics as external review, the Cluster employs philosophers, anthropologists and legal scholars who work alongside machine learning PhD students from project inception, enabling ethical considerations to shape research design.

Institutional funding should add to existing resources

An important factor in Cluster success is ensuring specialised funding represents a net increase in resources rather than substituting for existing budgets. Excellence funding should augment core institutional support, not displace it: if universities use Excellence grants to offset or withdraw departmental budgets, the catalytic effect is lost. Used appropriately, Cluster funding draws departments into investing additional resources of their own, as Tübingen's departmental AI hiring suggests is happening.

Technical infrastructure requires technical expertise

Funding for compute and hardware is often ample, but hardware requires people to maintain it. The difficulty of competing with private-sector salaries for technical experts such as GPU cluster managers suggests future funding models should include flexible salary brackets or dedicated technical talent funds.



AI FOR SCIENCE:

Case Study 6

Alchemy: the UK Hub for AI in Chemistry, led by Imperial College London and the University of Liverpool

A national hub to accelerate the development and use of AI in experimental and computational chemistry.

KEY TAKEAWAYS

- **Flexible funding broadens participation and keeps a hub connected to the frontier.** The Alchemy Hub's budget included a Flexible Fund of approx. £4.5m of unallocated funds, letting researchers across the network bid for funding as the hub develops. This openness draws early-career researchers into the ecosystem, encourages interdisciplinary projects that break down institutional silos, and lets funding follow the strongest ideas rather than choices fixed at proposal stage.
- **Research software engineers are essential translators.** Embedding professional software engineers within projects reduces friction between domain experts and AI specialists, and ensures complex research is translated into functional, scalable tools. The NMR Data Hub at Imperial – built by an Alchemy research software engineer – shows the kind of infrastructure that dedicated technical staff can deliver.
- **Publicly funded hubs give the private sector a neutral venue for cooperation.** The hub model lets companies come together and openly discuss what is needed for the sector to up-skill and take up the opportunities that AI offers. Through roundtable discussions with industry partners, the shared needs for digital literacy of chemists have been mapped and will be shared more widely with the chemistry community.

CONTEXT AND CHALLENGE

UK chemistry has had pockets of world-class capability in AI and automation, but no national coordination. Chemistry research is often siloed across institutions and sub-disciplines, and data infrastructure is fragmented. Experimental data has been scattered across local drives, lab notebooks and inconsistent formats, making it hard to train or apply machine learning models. Most researchers lack routine access to the AI and software engineering skills needed to use these techniques effectively, and there have been few structured opportunities for chemists and AI researchers to work together.

Imperial College London began building chemistry-automation capability in 2019 with ROAR (Rapid Online Analysis of Reactions), a high-throughput facility for automated chemical reactions, followed by DigiFAB (the Institute for Digital Molecular Design and Fabrication) and a Master's in Digital Chemistry. The University of Liverpool built complementary strength in materials discovery, anchored by the £81m Materials Innovation Factory, co-created with Unilever and opened in 2018, with their own Master's in Digital Chemistry following in 2024.

In February 2024, Alchemy received £12m from the Engineering and Physical Sciences Research Council (EPSRC), supplemented by £19m in cash and in-kind contributions from consortium partners, to combine these capabilities and coordinate across the AI and chemistry ecosystem. The consortium spans nine UK universities – Liverpool, Imperial, Queen's University Belfast, UCL, Cambridge, Manchester, Oxford, Southampton and Strathclyde – plus more than 15 industrial partners, with a remit running from molecules through materials to devices. Alchemy's challenge is to turn AI in chemistry from a specialism into routine practice: to give UK chemistry researchers the data, tools, infrastructure, skills and collaborators needed to apply AI methods, and to do so quickly enough to keep pace with the field.

APPROACH AND IMPLEMENTATION

Funding follows ideas

£4.5m of Alchemy's funding was not pre-allocated to named projects when the hub launched. Instead, researchers bid for funding through open calls. The first major call – the £2.5m Frontier Fund – drew more than 50 applications requesting £44.4m. A separate pump-priming call funded ten projects with £250,000 in total to seed early-stage work at the AI–chemistry interface. This model engages the wider community and lowers the barrier for early-career researchers, who would struggle to win named-investigator lines in a pre-allocated grant. It also lets the hub respond to ideas that did not exist when the original proposal was written.

Distributed infrastructure, shared skills

The hub is an equal partnership between Imperial and Liverpool. Research direction is set by domain leads drawn from across the consortium, informed by input from both senior and early-career researchers, and supported by project managers and research software engineers. Rather than building a single centralised facility, Alchemy researchers draw on existing physical infrastructure across the consortium: Imperial provides automation capability through DigiFAB, and Liverpool provides materials discovery infrastructure through the Materials Innovation Factory. Researchers have access to in-house high-performance computing and GPU resources, though demand from AI workloads now routinely outstrips supply.

Alchemy treats community engagement as core infrastructure. In 2025 the hub ran nine webinars, three training schools and workshops, and 12 live events including research conferences, reaching more than 1,500 attendees in total; 14 undergraduate internships were placed during the year. The Master's in Digital Chemistry at Imperial, launched ahead of the hub, trains chemists in data handling, programming, chemical automation and AI methods alongside core chemistry, with 2026 seeing the largest cohort of students showing increased demand for such courses.

Making chemistry data AI-ready

Effective AI in chemistry requires well-structured, accessible data. Alchemy's flagship response is the NMR Data Hub at Imperial. The Chemistry Department's nuclear magnetic resonance facility runs four Bruker instruments and is one of the most heavily used UK facilities for compound characterisation. An Alchemy-funded research software engineer rebuilt how its data is captured, stored and accessed: a web portal now stores each experiment with its associated metadata in line with FAIR data principles (Findable, Accessible, Interoperable, Reusable), and an API layer is being built on top so that data flows directly into machine learning pipelines rather than being downloaded and cleaned by hand. Implementation required replacing unsupported systems and updating software that had not been touched in years. The resource is open source, in line with Alchemy's commitment to open data and tooling, though hub members remain free to use commercial alternatives.

IMPACT AND OUTCOMES

In its first full operational year, Alchemy delivered:

- Ten pump-priming projects funded (£250,000 in total), seeding early-stage AI–chemistry research.
- Seven peer-reviewed publications in 2025.
- A Frontier Fund call drawing more than 50 applications requesting £44.4m against a £2.5m budget – a measure of latent demand in the community.
- More than 1,500 attendees across nine webinars, three training schools and 12 live events, and 14 internships placed.
- The NMR Data Hub at Imperial – the first deployment of a FAIR-aligned data infrastructure for an Alchemy-affiliated facility.

Beyond the hub's own funding period, Liverpool has announced longer-term plans for AIM-HI – the AI Materials Hub for Innovation – a £100m AI-driven materials project within a £111m purpose-built chemical sciences facility scheduled to open in 2031, intended to scale the Materials Innovation Factory's industry-collaboration model nationally.

LESSONS FOR POLICY

Education must span the professional spectrum

The chemist of the next decade will need data handling and AI skills alongside fundamental chemistry, this was a clear outcome from our industry roundtable on AI Skills for Chemistry. Existing routes – undergraduate, PhD, postdoctoral – were not designed for that combination. Imperial and Liverpool's Master's in Digital Chemistry are one model; Alchemy's training days, schools and internships are another. Capability building needs both dedicated degree-

level programmes that train chemists in AI from the start, and continuing-professional-development routes that retrain mid-career researchers.

Structural incentives must reward interdisciplinary team science

AI-driven chemistry depends on teams – chemists, AI researchers, materials scientists, software engineers, automation specialists. Alchemy has embedded an interdisciplinary team-based approach in its forerunner research projects, and those funded through the Flexible Fund. In addition, the hub-level governance spans multiple disciplines, hub researchers get shared credit on outputs. Such team-based science is not always evident across universities, but the Alchemy Hub is showcasing how this can work through incentivising structures that reward team science directly.

Research software engineers need career paths

The NMR Data Hub was built by Alchemy-funded research software engineers. Without those engineers the modernisation does not happen, and the data does not become usable for AI. That role – translating between chemists and AI methods, and building production-quality infrastructure – is essential to almost every project the hub funds. Yet universities struggle to retain research software engineers: salaries lag industry, contracts are short-term and project-tied, and promotion routes assume the traditional academic-publication model. Alchemy can fund these roles for the life of the grant; structured career paths are needed if hubs like Alchemy are to keep this talent in the UK academic ecosystem.

Agility is needed in a fast-moving field

By having a substantial Flexible Fund, Alchemy can move money toward the strongest ideas, regardless of who proposed them. Fixing project lists at proposal stage locks in choices made before the science is ready. Hubs that can re-allocate funding during their lifetime are structurally better placed to track the frontier.



AI FOR SCIENCE:

Case Study 7

The UK Centre for Ecology and Hydrology (UKCEH)

How an independent research organisation becomes AI-ready.

KEY TAKEAWAYS

- **Partnerships between universities and independent research organisation are mutually beneficial.** Universities translate cutting-edge methods, training capacity, and experimental space into education, while UKCEH translates it into national capability to benefit environmental sciences communities.
- **Making national data 'AI-ready' is an infrastructure investment.** UKCEH is transitioning decades of environmental monitoring data from legacy formats to machine-learning-compatible standards. This is time-consuming, technically unglamorous work that does not naturally attract funding. Policy needs to recognise data preparation as essential infrastructure.
- **The 'sufficiency' principle matters for responsible AI research.** UKCEH is developing guidelines that ask what level of compute is sufficient to achieve a scientific outcome, rather than defaulting to maximum capability, challenging the assumption that more compute is always better.

CONTEXT AND CHALLENGE

The UK Centre for Ecology & Hydrology (UKCEH) is an independent research institute that delivers environmental science and manages the UK's national capability in environmental monitoring, modelling and data preservation. UKCEH and its predecessors have modelled environmental change for over 60 years. Today, government departments, regulators, and researchers depend on its data and analysis to inform policy on land use, biodiversity, water quality, and climate adaptation.

Environmental data are growing in scale and complexity – higher-resolution sensors, continuous monitoring, real-time streams – and machine learning offers analytical capabilities that traditional physics-based models cannot match. Staying at the frontier requires UKCEH to become AI-capable.

This is not straightforward for an independent research organisation. Machine learning specialists command salaries that independent research organisations and public bodies cannot match; UKCEH's reputation depends on the trustworthiness of its outputs, so experimental AI tools cannot be adopted without validation; and an organisation dedicated to environmental protection cannot ignore the energy and resource costs of large-scale computation.

Environmental data is also heterogeneous, often incomplete, and structured around spatial and temporal scales that general-purpose machine learning tools are not designed for. Bridging the gap requires sustained collaboration between people who understand each side, which is why university partnerships, providing cutting-edge methods, training capacity and experimental space, have been central to UKCEH's transformation. Partnerships also create dependencies: when university researchers develop tools and move on, who maintains them? And how does an independent institute balance the openness universities expect with the governance requirements of nationally important data?

APPROACH AND IMPLEMENTATION

The DSNE partnership with Lancaster University

The Data Science of the Natural Environment (DSNE) project is a five-year, £2.5 million collaboration with Lancaster University. It was structured around a shared physical hub designed to foster the informal, cross-disciplinary communication needed to translate between statistical methods and environmental problems. UKCEH contributed decades of environmental data, and both institutions contributed statistical and AI expertise and computational capability.

The project developed new techniques tailored to environmental data, including advances in changepoint analysis for detecting shifts in ecosystems, and extreme value theory for modelling high-impact but data-scarce events such as floods and droughts. Experimental work was carried out on the DataLabs platform, which lets users of varying expertise analyse and visualise environmental datasets in ways that are open, transparent, and reproducible. Most importantly, DSNE demonstrated the value of sustained interaction between environmental scientists and data specialists, and led directly to the establishment of the Centre of Excellence in Environmental Data Science (CEEDS).

CEEDS and internal capability

CEEDS is a transdisciplinary community of over 350 environmental scientists, statisticians, software engineers, and domain experts. Recognising that the environmental sector cannot compete with technology companies on salary, it pursues internal upskilling rather than recruitment at scale. This includes a seminar series, spanning topics from machine learning for sensor data quality control to environmental digital twins, and a community of practice in which researchers share methods across projects rather than maintaining siloed expertise.

UKCEH also contributes to a shared ecosystem of tools rather than building everything internally. The European-funded Insect AI network shares species-identification classifiers for bees, butterflies and moths across institutions, achieving taxonomic scale no single institution could maintain, while UKCEH's Data Science Toolbox – a curated catalogue of methods as Jupyter Notebooks – is openly available for others to adapt and reuse.

Making data AI-ready

At the heart of UKCEH's transformation is the challenge of making decades of environmental monitoring data usable for machine learning. Much of this data has been stored in legacy formats such as CSV files and spreadsheets, while modern machine learning requires high-performance formats and consistent structure. In response, UKCEH is evolving its architectures to support standards such as Parquet and Croissant. The Environmental Information Data Centre is central to this effort. Beyond long-term preservation of nationally important datasets, a dedicated team of data stewards ensures all assets meet FAIR principles (Findable, Accessible, Interoperable, Reusable); if the underlying data are poorly documented or inconsistently structured, downstream models cannot be trusted.

Discoverability remains a persistent challenge. Creating metadata typically falls to scientists who see it as an administrative burden. UKCEH is developing tools that use large language models to extract metadata from existing documentation, and exploring natural language interfaces that would let users request, for example, 'all data related to health outcomes in the Lake District National Park'. These tools include built-in protections against hallucination: AI assists with extraction and discovery, but scientific integrity checks remain in place, implemented by human reviewers.

AI in the science: biodiversity monitoring at scale

UKCEH is exploring how AI methods can supplement – or in some cases emulate – the physics-based models that have historically underpinned environmental modelling. Real-time monitoring has also made manual oversight of data quality impractical, so machine learning workflows now identify anomalies and automate quality assurance at scale.

Biodiversity monitoring illustrates the transformation most clearly. Classifying species abundance has traditionally required labour-intensive field surveys. UKCEH now deploys remote cameras and audio recorders, training AI systems to classify species from visual and acoustic signatures. A single pilot has expanded to 180 active deployments, with international partners deploying the systems globally, enabling continuous, large-scale biodiversity assessment while freeing expert ecologists for analysis rather than data collection.

Compute: from internal infrastructure to hybrid partnerships

UKCEH historically maintained its own high-performance computing, but the shift from CPU- to GPU-based computation, combined with global hardware scarcity, has made a purely in-house model difficult to sustain. The organisation is moving to a hybrid strategy: some internal systems, partnerships with

national HPC centres, and a partnership with Amazon Web Services for cloud AI services difficult to replicate locally. The strategic risk is mitigated through a cloud-native principle: software is developed to be cloud-independent, so tools can be redeployed across providers or moved back in-house if pricing or access conditions change. This is treated as a safeguard for sovereignty over national environmental monitoring, reconciling pragmatic use of commercial cloud AI with the sufficiency principle below: exploit what is useful, but never in a way that cannot be unwound.

A 'human-in-the-loop' mandate applies throughout this infrastructure. Code generation is never fully delegated to machines, but is overseen by trained software engineers to ensure quality and transparency, as well as scientists to ensure it suits their needs.

Responsible innovation and the sufficiency principle

For an environmental organisation, the energy and resource costs of computation directly contradict its mission if not managed carefully. To approach this challenge, UKCEH convened a symposium of 60 experts to establish principles for its ethical framework, such as transparency and keeping scientists in the loop. A key concept emerging from follow-up research from the sustainable Digital Research Infrastructure (sDRI) is sufficiency: rather than defaulting to the largest models or maximum compute, guidelines ask what level of computational resource is sufficient to achieve the desired scientific outcome. This challenges the prevailing assumption that more compute is always better. For an institute accountable for both environmental impact and public funds, it is also a practical operational principle.

Data access and the role of an independent research institute

As an independent research organisation managing datasets funded by UKRI, the Environment Agency and others, UKCEH has a clear mandate to keep its outputs accessible: data produced through European or UKRI funding must remain open, though a small number are licensed and this proportion may grow under commercialisation pressures. Data silos are a persistent problem: environmental policy increasingly requires integration across climate, health, and social science data managed by different organisations with different access policies and technical standards. UKCEH's discovery tools make disparate resources easier to find, but this addresses symptoms rather than the underlying fragmentation.

IMPACT AND OUTCOMES

DSNE's methodological advances are now used in environmental research beyond UKCEH, and the project stands as proof of concept for university–public body partnerships in data-intensive environmental science. Its most consequential outcome, though, was institutional: the cultural shift that produced CEEDS.

That capability has translated into operational change: biodiversity monitoring has moved from manual field surveys to automated AI-driven systems at 180 deployments, and the openly available Data Science Toolbox has enabled external researchers to adapt UKCEH's tools for their own work.

LESSONS FOR POLICY

Partnerships are essential to help public research bodies build AI capability

Universities provide cutting-edge methods, training capacity, and experimental space. Partnerships work best when there is clarity about what each partner contributes, and about who is accountable when tools transition from experimental to operational use.

Making national data AI-ready is infrastructure investment

Transitioning decades of monitoring data to machine-learning-compatible standards is time-consuming, technically unglamorous, and does not produce publishable outputs; research grants fund new methods, not the painstaking work of standardising historical data. If the UK's national data assets are to be AI-ready, this needs to be treated as a standing infrastructure cost.

The sufficiency principle challenges default assumptions about compute

UKCEH's sufficiency guidelines are a practical response to its dual accountability for environmental impact and public spending. Much AI policy assumes that more compute is better and that the constraint is access. UKCEH's approach suggests a different question: when is additional compute necessary, and when is it wasteful?

Research institutions deploying AI need to consider governance and accountability for operational science

When UKCEH develops biodiversity monitoring tools, the question is not only whether they work scientifically but whether regulators can rely on them operationally. The 'scientist-in-the-loop' principle reflects this. Funding systems support university research and government procurement of services; the intermediate space, where public research bodies translate experimental tools into operational infrastructure, fits neatly into neither. Clarifying how this translation is funded and governed would strengthen the pipeline from research to policy impact.



AI FOR SCIENCE:

Case Study 8

European Centre for Medium-Range Weather Forecasts (ECMWF)

How a 50-year-old intergovernmental physics institute put a machine learning model into operational forecasting.

KEY TAKEAWAYS

- **Build trust by making machine learning outputs contestable.** ECMWF publishes forecasts from its Artificial Intelligence Forecasting System (AIFS) on the same public charts as its physics model, and gives Member States the data to stress-test the model on their own terms. Documenting known limitations and letting users find the failure modes builds trust.
- **Enable close working between machine learning expertise and scientific teams.** AIFS was built by mixed teams of atmospheric physicists and machine learning scientists, bringing physics knowledge into model design and heading off a common failure mode: a new AI team running in parallel to, and in tension with, the scientific core.
- **Fund open frameworks as well as open products.** One of the most durable outputs of ECMWF's work is Anemoi, an open-source framework co-developed with eleven European national weather services, which lets each country train models on its own data while sharing tooling.

CONTEXT AND CHALLENGE

ECMWF is an intergovernmental organisation founded in 1975 and primarily funded by 23 Member States and 12 Co-operating States. It is both a research institute and a 24/7 operational service producing global weather forecasts used by national meteorological services across Europe and beyond. For 50 years its forecasts have been generated by a physics-based model, the Integrated Forecasting System (IFS), which solves the equations of atmospheric motion on a global grid and is widely regarded as the most accurate global weather model in operation.

Physics-based forecasting is expensive. A single global run consumes substantial supercomputer time and electricity, limiting how often forecasts can be refreshed. ECMWF researchers began testing machine learning in the 2010s, and the shift accelerated after the 2018 release of ERA5: a consistent hourly reconstruction of global weather since 1940 that provides the large, clean, labelled dataset machine learning needs. By 2022, external groups – including Huawei, NVIDIA, and Google DeepMind – had shown that models trained on ERA5 could match or beat physics-based forecasts on several measures, at a fraction of the compute cost.

The question for ECMWF was no longer whether machine learning worked, but how an established scientific institution – with 550 staff, a dual research-and-operations mandate, and 35 governments watching – should integrate it without destabilising a service governments and industry already depend on.

APPROACH AND IMPLEMENTATION

From experimental model to operational system

ECMWF built its own machine learning forecasting system, the Artificial Intelligence Forecasting System (AIFS), by first running open-source academic models internally to understand their behaviour, then building their own end-to-end system. AIFS spent a period as an experimental model – published openly, fed live data, but flagged as not for operational use – while Member States and users stress-tested it. In February 2025, AIFS Single (a deterministic 28km-resolution model, running a single forecast at a time) became operational; in July 2025, AIFS ENS (a probabilistic ensemble of 51 forecasts at 31km resolution) followed.

'Operational' means AIFS forecasts are produced on the same fixed schedule as the IFS, disseminated through the same pipelines, and published alongside IFS outputs on ECMWF's OpenCharts. The two systems use the same observational inputs and are verified against the same ground truth; Member States choose which products to use and for what purpose.

The experimental phase surfaced failure modes that technical benchmarks had not. This included occasional atmospheric artefacts in temperature or pressure fields that are statistically plausible but physically odd. Every AIFS product was labelled experimental, known limitations were documented in user-facing notes, and Member States were advised which decisions the model should and should not yet be used for.

Distributing machine learning expertise across teams

ECMWF did not create a separate AI department. Machine learning specialists were embedded into existing scientific teams alongside atmospheric physicists, and mixed teams built AIFS together. This cultural choice forced physics knowledge into the model design and prevented an AI unit and the scientific core drifting into competition for resources and credibility.

The distributed model is supported by internal training: in-person courses at ECMWF's Reading headquarters for staff and Member State forecasters, an EU-funded series for meteorologists who are not machine learning specialists, and a massive open online course launched in 2023 that has attracted more than 9,000 participants from 159 countries.

Building shared infrastructure, as well as a shared model

AIFS is open source. The framework used to build it – Anemoi – was also released in 2024 as a joint project between ECMWF and eleven European national meteorological services, including Spain's AEMET, Germany's DWD and Météo-France. Anemoi is a Python toolkit handling the full pipeline – preparing training data, defining model architectures, training, and running the resulting model operationally – and now underpins AIFS as well as several national systems, including Germany's AICON.

A related project, Forecast in a Box, packages Anemoi-built models into reproducible pipelines that run on modest hardware or in the cloud, for forecasters without access to a supercomputer. Sharing the framework multiplies capability across the network.

Verification: adapting existing systems rather than building new ones

ECMWF tests AIFS through the same verification infrastructure it uses for IFS. The model is run daily, scored against observations, and published in standard forecast skill scores; no new regulatory framework was needed. The adaptation required was specific: because machine learning models are trained to optimise a cost function, they can score well on headline metrics while generating subtly non-physical behaviour a physics model would not produce. ECMWF's verification teams are therefore working with Member State services to widen the metrics tracked, catching blind spots the legacy verification never needed to check for.

Hardware: buying flexibility, not just compute

ECMWF's own supercomputers are CPU-heavy, configured for physics simulations rather than GPU-intensive model training. Instead of replacing its hardware, ECMWF accesses GPU capacity through EU-funded EuroHPC facilities such as LUMI in Finland and MareNostrum 5 in Spain. It also keeps its software stack deliberately portable – built on PyTorch, with data in open formats – so models run on NVIDIA, AMD or other hardware without rewriting. This insulates the Centre from vendor lock-in and single-site outages.

Keeping the archive sovereign while making the data open

ECMWF's machine learning work depends on ERA5 and decades of observational data held on a physical tape archive in its Bologna data centre. Tape is slower than cloud storage but cheaper at scale, and keeps the primary asset under direct institutional control. Rather than migrating the archive to a commercial cloud, ECMWF partners with organisations that host clones of selected datasets for fast researcher access, and licenses the data permissively through the EU's Copernicus programme so that models trained on it can be commercialised and redistributed freely. The archive and the open mirrors do different jobs: one preserves institutional control, the other removes friction for users. 'Open' and 'sovereign' are design choices that can coexist if the architecture is planned for it.

IMPACT AND OUTCOMES

On headline measures, AIFS is competitive with or better than IFS. For tropical cyclone tracking, AIFS outperforms state-of-the-art physics models by up to 20% – roughly 12 hours of additional lead time on a storm's path – with similar gains for day-5 surface temperature in the ensemble version. AIFS produces forecasts more than ten times faster than IFS and uses roughly 1,000 times less energy per forecast.

These results come with limits. AIFS runs at coarser resolution than IFS (28km against 9km), so IFS remains indispensable for fine-scale features such as localised precipitation. Machine learning models also degrade differently from physics models. Trained on physical observations, they are bounded by the structure of the data in a way language models are not – the probability space of tomorrow's atmosphere is far more constrained than the probability space of the next word; AIFS does not fabricate continents – but they can produce subtle physical inconsistencies, such as smoothed extremes, that a physics model would not. The governance question is not 'will it hallucinate' but 'which variables, previously guaranteed by physics, now need explicit monitoring'. Parallel operation matters partly because the two systems fail in different ways, making each other's blind spots visible.

Financial and energy firms – which care about speed, tolerate coarser resolution, and already use multiple models – moved quickly to make use of the new capabilities. Public-safety use, such as triggering evacuation warnings, is slower, because the track record required before a national agency will stake lives on a new model is longer than any machine learning system has yet accumulated. Parallel operation is what lets that track record build.

LESSONS FOR POLICY

Make machine learning outputs contestable, and adapt existing verification

Publishing AIFS alongside IFS on identical inputs, for long enough to build a shared evidence base, is a strategy transferable to any scientific domain with an established operational model. It also shows a pathway for governance efforts: mature fields already have decades of verification infrastructure, and ECMWF validated AIFS through the same pipelines as IFS with targeted additions for machine-learning-specific failure modes. Policymakers should fund the adaptation of existing verification systems where mature scientific verification exists; in fields that lack such infrastructure, such as much of AI-for-biology, building it is a precondition for trustworthy deployment.

Embed machine learning expertise in scientific teams, not a parallel unit

ECMWF's decision not to create a separate AI department – embedding machine learning specialists in existing scientific teams instead – brought physics knowledge into model design and avoided the competition for resources and credibility that a stand-alone AI unit invites. Internal training is what makes the distributed model work. It gives existing staff a route into the technology rather than positioning AI as something done to them.

Fund open frameworks as well as open products

Anemoi – co-developed with eleven national weather services and now underpinning several national forecasting systems – shows that sharing a framework multiplies capability across a network in a way that sharing a finished model does not. For a funder, financing the framework produces more capability than financing individual models.

Institutional funding structure shapes what is possible

ECMWF's transition was enabled by a budget structure most national scientific institutions do not have: annual contributions from 35 governments on multi-year cycles, a formal dual research-and-operations mandate, and the authority to reallocate staff to new priorities without cutting core services. This let ECMWF treat AI integration as a complement to IFS rather than a competitive threat. For example, it could hire external talent while also opening space for existing staff to engage with the technology, supported by the psychological and financial safety to experiment. An organisation's ability to pivot toward AI depends on funding that allows exploratory activity without cannibalising core services; where budgets force a choice, the exploration loses.



AI FOR SCIENCE:

Case Study 9

The OpenBind consortium

A pioneering initiative creating the world's largest open-access dataset of drug-protein interactions.

KEY TAKEAWAYS

- **Scaling AI data generation requires clear and reproducible methods.** The ambition to generate data at scale cannot rest with one team. Instead, it should become a way of working so that shared processes are used and all experiments can contribute to the goal. OpenBind's consortium structure and culture ensure that this shared standard of scientific rigour flows across its network.
- **Data initiatives must adopt a production mindset from day one.** Complex research projects require regular meetings and continuous action-tracking to ensure they move at speed, and send a signal that the endeavour is serious. OpenBind treated its setup as a production problem from the outset, establishing immediate weekly meetings with minutes, templates and actions. This allowed funding to be deployed rapidly and effectively, and set an ambitious pace.
- **Ambitious scientific endeavours require courage.** Committing to a project of OpenBind's magnitude requires an anchor institution willing to demonstrate the courage to execute a new project at speed, plus strong governance and organisation. OpenBind's lead anchor institution - Diamond Light Source - has a track-record in leading ambitious projects and an understanding that strong governance is required if funding must be spent rapidly.
- **There are multiple policy lessons.** AI data generation must be funded at scale and over the long term; it requires dedicated strategic funding calls; operational lessons can be applied across scientific domains; and data sovereignty should be recognised as a national asset.

CONTEXT AND CHALLENGE

AI models are only as good as the data they are trained on. Yet, the data required to build AI tools for discovering new medicines and advancing science is surprisingly scarce in the public domain.

OpenBind aims to address this barrier by building the world's largest open-access dataset of protein-ligand interactions, as a foundational national capability aligned with the UK Government's AI for Science Strategy and the UKRI AI Research and Innovation Strategic Framework. Rather than focussing on a particular challenge or targeting isolated drug pipelines, OpenBind is generating clean, scalable FAIR datasets for the wider scientific community - using automated chemistry and high-throughput X-ray crystallography - and running regular blind prediction challenges. This will unlock AI models that generalise across targets and thus diseases, and specifically better and knowably reliable predictions of how molecules bind to these targets; and ultimately, this should tremendously speed up drug discovery, from target validation to lead optimisation. In this way, it is not just a dataset project, but a continuously operating discovery platform linking experimental data production to real drug-design programmes.

OpenBind is transitioning from its startup phase, moving from establishing a robust proof of concept: it has now developed the processes, protocols, and operational foundations needed to generate high-quality, AI-ready structural data at scale, and is now scaling and industrialising delivery and thereby turning the UK's AI for Science ambitions into sustained, long-term capability.

APPROACH AND IMPLEMENTATION

Consortium governance and funding

OpenBind was launched with £8 million in seed funding from DSIT's Sovereign AI Unit, structured as a multi-institutional consortium coordinated by Diamond Light Source, which acts as the central hub for the initiative. Academic leadership comes from the University of Oxford alongside global partners, combining national research infrastructure, higher education, and industry.

Entrusting the project to a world-leading facility with a strong reputation helped secure government funding. Yet, taking on rapid-expenditure commitments required significant institutional courage from Diamond Light Source's leadership, who swiftly established strong governance processes enabling funding to be spent quickly and effectively.

An enormous number of skills are required to work on the multidisciplinary challenge of generating AI-ready data. The consortium structure leverages domain expertise across partners while protecting against single-point-of-failure risks, creating a consensus of rigour with all parties working towards a common goal to solve the problem defined at the outset. Articulating this consensus properly was key to smoothing the path to legal agreements, which is invariably a slow step in all formal consortia. The longevity of the project has recently been ensured with two awards: a $2.87m Gates Foundation grant to explore how OpenBind's fragment-based drug discovery approach can be applied to urgent global health challenges; and a £3.4m Wellcome grant that partners OpenBind with BioStruct-Africa, to develop equitable access for AI-based drug discovery. Its overall funding ambition for the next 2 years is £10-20m.

Leveraging national infrastructure

OpenBind's operational model builds on its experimental facilities that already run at scale. The project relies on Diamond Light Source's world-leading high-throughput X-ray crystallography facilities, which are now augmented with cutting-edge automated chemistry. The aim is to spearhead a global drive to generate more than 500,000 protein-ligand structures over a period of five years, which increases by twenty-fold what came into the public domain over the previous 50 years. This experimental output is paired directly with national compute assets, such as the Isambard AI supercomputer facility in Bristol, to train and validate predictive models against incoming experimental results. Placing the initiative within proven, stable infrastructure ensures the operational consistency needed to generate large-scale, FAIR datasets.

Drawing on lessons from the COVID Moonshot, OpenBind adopted a production mindset from day one, establishing centralised document sharing and an intense meeting schedule to track complex pipelines. Trust among key leads allows responsibilities to be shared effectively across institutions.

Enabling cross-consortium collaboration

To manage hundreds of interconnected tasks spread across multiple meetings and a complex set of experiments required for data generation for data generation, the consortium relies on regular meeting cadences and clear action-tracking to keep the project plan in place. Crucially, this operational rhythm has driven a cultural shift among experimentalists, reframing data plus methodology and capture information as ongoing assets.

From the start, OpenBind worked hard to generate a collective understanding of the ambitions, methodologies and problems. The main mechanism was fortnightly science seminars for the wider team, along with staff exchanges and a small number of in-person meetings. Even so, the complexity and diversity of the project, and its rapid pace, meant it took the better part of a year to approach this goal; and in retrospect, the knowledge exchange should have had a bigger allocation of time and cost.

What did work and was essential, were strict meeting cadences during initial setup, to manage operational complexity; but this meeting load was neither sustainable nor necessary longer-term, so the operational model evolved as needed while ensuring effective communication. As the project matures, OpenBind aims to ease central orchestration, eventually enabling partner institutions and beamlines globally to operate independently while keeping the data flowing.

Ensuring transparency and data stewardship

Open access data and transparency are at heart of OpenBind's project. While AI claims from closed models (typically commercial) are hard to assess and readily inflated, OpenBind places importance on data stewardship and blind evaluation in order to check and validate claims. Commercial drug discovery typically embargoes data for intellectual property reasons; instead, OpenBind embargoes its data for scientific reasons, namely to achieve robust validation of algorithms. Eventually, though, these high-quality protein-ligand datasets are released unencumbered, ensuring that accurate training data is accessible to the wider scientific community, and also available for independent verification. The consortium believes it has a duty to use scientific blinding to validate results and hopes other projects will benefit from its rigorous project.

Building industry partnerships and navigating commercial barriers

OpenBind is now actively building an industrial consortium within which commercial partners help steer the generation of useful experimental AI-ready data, in turn providing OpenBind with the commercial expertise and real-world use cases needed to ensure its outputs are aligned with real-world drug discovery. Many of OpenBind's delivery partners had already worked together in previous consortia, so trust and credibility were already established. In contrast, securing commercial engagement for the industry consortium has presented the expected roadbumps, not least that businesses underestimate the cost of generating clean, well-structured data, or conversely cannot accommodate the long-term nature of the business proposition. In general, only a fraction of potential partners can approve this internally, which means the net must be cast widely; conversely, those that do sign up are then highly committed.

A secondary challenge lies in aligning commercial expectations of data embargo with the rapid pace of the project. Traditional industry standards favour delays of 2-3 years, but a shorter timeline is needed to work within the context of blind challenges and regular data generation in a fast-moving field. This tension necessarily requires compromiseand, to reach that, the commitment of the partners is an essential ingredient. Such a balance is however possible: for instance, the embargo period required for commercial discovery can be aligned with a high cadence of data release, by generating data that will drive the model to generalise, while keeping it adjacent to commercial discovery as opposed to overlapping.

Ultimately, what this needs is a collaborative approach so that orthogonal goals support one another, and are not considered as conflicting priorities. Thus OpenBind provides scientific leadership and generates academic outputs, but it is the engagement with industry that ensures its work actually can contribute to commercial discovery pipelines, as well as aligns with government goals. Consortium partners thus join OpenBind not only to meet their own objectives but because they are community minded. With this spirit of cooperation and a well-defined shared aim and way of working, publications effectively write themselves and shared aims are more easily met.

IMPACT AND OUTCOMES

OpenBind (OpenBind, 2026a) delivers value across multiple dimensions of the drug discovery ecosystem. With its first data release in May 2026, comprising publicly available high quality, standardised experimental data and a newly trained predictive model, the consortium demonstrated that its pipeline is operational, having generated 800 high-quality measurements in only 7 months. In the past, such a large dataset would take years to be produced and would almost never be released as a unified entity.

The subsequent release in August 2026 of OpenBind-0 (OB0) (OpenBind, 2026b), marks the first in a new series of fully open-source molecular structure prediction co-folding models built on OpenFold3 and specialised for predicting the structures of proteins bound to small-molecule ligands. The consortium also released a dense, discovery-relevant dataset of 717 ligand-bound structures that capture fragment-to-hit progression across three drug targets.

Finally, OpenBind announced in October 2026 its first blind challenge (OpenBind, 2026c)., in collaboration with the OpenADMET consortium, which will illuminate the state-of-the-art of cofolding models for real-life drug discovery problems, as opposed to the retrospective benchmarks that have been the norm in the field to date.

At a lower methodological level, the consortium is developing an ontology about how affinities can be communicated, which is a huge gap in the field: unlike for protein structure, there is very little standardisation of the deposition of ligand-protein affinity data.

At the cultural level, OpenBind bridges a gulf of understanding between experimentalists and computational modelers by facilitating a constant conversation that has proved to be insightful. By ensuring that experimentalists are part of the model design and understand what the models are doing, the project dismantles silos and creates a single operation enabling collaboration and constructive conversations that enable progress.

LESSONS FOR POLICY

Fund AI data generation at scale and over the long-term

Models require large, purpose-built datasets. Creating them is expensive and resource-intensive, requiring a large upfront investment. While seed funding enables ambition, long-term, large-scale funding is required with enough freedom to prevent the risk of killing initiatives before consortia can be set up, work effectively and generate value.

Shift from reactive funding calls to proactive strategic ideas

Current funding landscapes tend to favour project calls rather than issuing bold challenges that force communities to assemble around strategic ideas. OpenBind succeeded because it was already considering how to generate useful data, while others were unprepared to secure funding.

Replicate operational success across scientific domains

OpenBind offers a blueprint that could be replicated across numerous scientific disciplines when government funding is available, so that new projects can benefit from lessons already learned.

Recognise data sovereign capability as a national asset

AI tools require good quality structured data, making it a valuable commodity. Data can therefore be a national differentiator and an asset that will attract commercial investment and talent.

References

OpenBind (2026a) 'OpenBind's first data and model release marks a milestone for AI-enabled drug discovery', OpenBind, 5 May. Available at: https://openbind.uk/news/openbinds-first-data-and-model-release-marks-a-milestone-for-ai-enabled-drug-discovery/ (Accessed: 6 November 2026).

OpenBind (2026b) 'Blog: OpenBind-0: advancing open molecular structure prediction', OpenBind, 11 August. Available at: https://openbind.uk/news/blog-openbind-0-advancing-open-molecular-structure-prediction/ (Accessed: 6 November 2026).

OpenBind (2026c) 'Announcing the first OpenBind blind challenge: Zika virus NS2B-NS3 protease', OpenBind, 14 October. Available at: https://openbind.uk/news/announcing-the-first-openbind-blind-challenge-zika-virus-ns2b-ns3-protease (Accessed: 6 November 2026).



AI FOR SCIENCE:

Case Study 10

Polymathic AI

Building foundation models for science through international collaboration.

KEY TAKEAWAYS

- **Open data release can be more transformative than the models themselves.** Polymathic AI released 115TB of formatted scientific data before releasing its models, and external researchers used these datasets to build their own foundation models. Policy interventions focused on standardising and releasing scientific datasets may enable as much innovation as funding model development directly.
- **International collaboration at this scale depends on compute access from multiple jurisdictions.** Polymathic AI draws on US foundation funding, French national infrastructure and commercial NVIDIA allocations. Securing large allocations was difficult and managing access across systems created friction, suggesting a gap in international research infrastructure coordination.
- **Foundation models are computationally intensive to train but designed to reduce compute requirements at the point of use.** Whether this democratises access depends on governance: who hosts them, who controls access, and whether fine-tuning remains within reach of researchers with limited resources.

CONTEXT AND CHALLENGE

Foundation models – systems trained on large, diverse datasets that can be fine-tuned for specific tasks rather than built from scratch – have dominated recent advances in language, vision and multimodal AI. Fine-tuning a pre-trained model is typically more effective and efficient than training a specialist model for each new task.

Scientific machine learning has so far not followed this trend. Most applications remain domain-specific – AlphaFold for protein folding, specialist models for climate prediction, bespoke systems for particular physics simulations – delivering impressive results within their domains but unable to transfer knowledge across them. A model trained on astrophysics data cannot contribute to fluid dynamics research, even when the underlying physics shares common structure.

The challenge is representational. Foundation models for language or vision can learn from diverse data because that data shares a common format: text is text, pixels are pixels. Scientific datasets – observations, simulation outputs, experimental measurements – are fundamentally heterogeneous, with no unified numerical representation.

Polymathic AI was established in 2023 to solve this problem: building foundation models that learn transferable structure across scientific domains. The collaboration brings together researchers from the Simons Foundation and its Flatiron Institute, New York University, the University of Cambridge, Princeton University and the Lawrence Berkeley National Laboratory, working across physics, astrophysics, mathematics, AI and neuroscience.

APPROACH AND IMPLEMENTATION

Governance and collaboration

Polymathic AI has kept its governance lean to preserve agility. A Strategic Planning Committee provides high-level direction, but most technical and strategic evolution is driven by a core team of approximately 20 researchers, with consensus-driven decision-making centred on weekly strategy meetings. Compute and data priorities are settled through internal negotiation, but there has always been a unified stance on open-source distribution, which is reinforced by institutional mandates, since funding from the Simons Foundation and Schmidt Sciences requires or heavily encourages free release of all models. Polymathic AI has also benefitted from open-source tools and wants to contribute back to the community.

Intellectual collaboration across the US, UK, and France is not hampered by geography, but the international structure introduces operational bottlenecks around compute access.

Data: formatting and release

The first year of the project was spent collecting and formatting data. Unlike language models built by scraping the internet at scale, scientific data requires careful curation. The team compiled datasets from astrophysics, biology, acoustics, chemistry and fluid dynamics; much was already public, and simulation data was generally available once the team engaged the researchers who created it.

The primary challenge was converting diverse numerical datasets into a uniform structure suitable for machine learning. This was time-consuming work in which the team manually verified that physical variables such as density and temperature were correctly labelled and consistent. This bottleneck persists despite advances in AI coding agents: in the team's experience such agents take shortcuts and overlook errors, making them unsuitable for high-fidelity scientific datasets.

The team released two open-source dataset collections in 2024, totalling 115 terabytes. The Multimodal Universe contains hundreds of millions of astronomical observations (over 100TB) from sources including the Sloan Digital Sky Survey, Gaia, and the James Webb Space Telescope; The Well contains 15TB of simulation data spanning 16 scenarios, from merging neutron stars to atmospheric dynamics. Both are freely available on HuggingFace and other repositories. Because they were released before Polymathic AI's own models, external researchers were able to build and release their own foundation models trained on them. The datasets have become a foundational resource for the field.

Models: Walrus and AION

Polymathic AI has developed two foundation models. AION was trained on the Multimodal Universe and works across different astronomical instruments and observation types; Walrus was trained on The Well and applies structure learned from one physical system to entirely different ones, handling problems as different as exploding stars, Wi-Fi signal propagation and bacterial movement.

By learning shared physical structure, the models perform better on new tasks than models trained from scratch, and with less data. This is the core value proposition of foundation models, and one that matters most precisely where data is scarce and compute is limited.

Compute and infrastructure dependencies

Compute comes from three main sources: the Flatiron Institute, NVIDIA's AI research resource (used for the physics model), and Jean Zay in France (used for the astrophysics model) operated by the French National Centre for Scientific Research (CNRS). Securing large allocations was challenging and there is no centralised compute budget: individual model leads secure resources within their regions. This diversifies dependencies, making the project less vulnerable to any single policy shift, but harder to replicate, because success partly depends on the personal networks of senior team members.

The international structure also creates day-to-day coordination challenges: user numbers on the Flatiron cluster are limited, and administrative barriers can delay access for UK-based researchers, which might sometimes be two months for a single student. The team trains on NVIDIA chips but otherwise uses open tools such as PyTorch. Dependence on foundation funding and national allocations creates a specific vulnerability: the team could not afford commercial compute rates if allocations were withdrawn. There are no signs of this happening, but there is a structural dependence.

Funding model

Polymathic AI operates through a hybrid, distributed funding model: Schmidt Sciences funds roughly five to eight key positions including postdocs and research engineers, while most researchers – particularly PhD students who align their thesis work with the collaboration – are supported by external sources.

Because the team relies on competitive in-kind grants rather than cash for compute, the ability to apply across multiple jurisdictions provides a safety net. When its application for compute on the UK's Isambard-AI supercomputer was initially rejected, it secured the necessary capacity through the French government instead, avoiding a regional bottleneck that might stall a more locally dependent initiative. Polymathic AI has since

been awarded gateway allocations on Isambard-AI.

Philanthropic organisations were willing to fund large-scale scientific foundation models – and to engage with the high-risk early stages – well before traditional government agencies. Philanthropy can provide the seed infrastructure for ambitious AI projects, which national resources can supplement once the technical path is proven.

IMPACT AND OUTCOMES

The most notable result to date is the catalytic effect of the data: Six foundation-model families now have documented training or evaluation lineage on The Well before Polymathic's own model releases. From a policy perspective, this suggests investment in data infrastructure may be more catalytic for innovation than investment in model development.

The models themselves demonstrate that the technical bet pays off: cross-domain transfer works, predictions remain coherent over time, and performance on downstream tasks exceeds prior foundation models. This allows faster adaptation to new problems and better performance in low-data regimes. Quantitative outcomes include the two dataset releases, the two open-source foundation models – Walrus (McCabe et al., 2025) and AION-1 (Parker et al., 2026) - and an influential 2023 paper on multiple-physics pretraining that shaped the field (McCabe et al., 2023). The Well has accumulated over 4,400 stars on GitHub, and the models are enabling research in domains where traditional simulations struggle, including projects investigating carbon dioxide removal in the ocean. As of September 2026, Polymathic outputs have been used in 184 studies and integrated in 12 pieces of software. Furthermore, Google DeepMind built upon Polymathic AI's open model by fine tuning Walrus's openly released weights for its Kastor tool (Couairon et al., 2026). The broader impact on research culture – whether foundation models become standard scientific tools – is still unfolding.

LESSONS FOR POLICY

Open data infrastructure may be as valuable as model development

By prioritising curation, standardisation, and open release, the project enabled a broader ecosystem of model development than it could have achieved by focusing resources on its own models. Open release does, however, leave a question about value capture: it does not guarantee that value accrues where the investment was made, or that essential tools remain governed in the public interest if they become commercially valuable.

International collaboration requires international compute coordination

This Cambridge-rooted, UK-linked collaboration at the frontier of scientific AI had to source its compute abroad. Polymathic AI now draws on compute from multiple national jurisdictions, found securing large allocations difficult, and experienced persistent friction managing access across systems. There is no straightforward mechanism for researchers to access compute across national boundaries, and coordination between systems is ad hoc.

Foundation models create governance questions as well as technical ones

Foundation models for science are intended to become shared research infrastructure: resources many researchers depend on, rather than tools built for individual projects. This creates governance questions single-purpose models do not: who hosts the models, who decides access terms, and – if a model becomes essential to a field – what ensures it remains available, maintained, and governed in the research community's interest? Open-source release addresses access in the immediate term, but it is not by itself a governance model: it does not specify who maintains models or what happens if hosting becomes expensive.

The two-tier risk is real but not inevitable

There is a concern that AI for science could create a two-tier system: institutions with access to compute and models, and those without. Polymathic AI engages with this by releasing models openly – lowering the barrier for researchers who cannot train foundation models themselves – and by designing models that reduce compute requirements at the point of use. But running foundation models still requires non-trivial compute, and fine-tuning requires expertise: researchers without GPU clusters or AI engineering support will still be at a disadvantage. Policy will need to address not just whether models are open, but whether researchers can use them.

References

Couairon, G., Jacq, A., Wu, Y.H., Singh, R., Hasson, Y., Berthet, Q. and Elie, R. (2026) 'Kastor: an efficient fine-tuning strategy for generative emulation of PDE simulations', arXiv, arXiv:2608.06107. Available at: https://arxiv.org/abs/2608.06107 (Accessed: 2 November 2026).

McCabe, M., Blancard, B.R.S., Parker, L.H., Ohana, R., Cranmer, M., Bietti, A., Eickenberg, M., Golkar, S., Krawezik, G., Lanusse, F. and Pettee, M. (2023) 'Multiple physics pretraining for physical surrogate models', arXiv, arXiv:2310.02994. Available at: https://arxiv.org/abs/2310.02994 (Accessed: 2 November 2026).

McCabe, M., Mukhopadhyay, P., Marwah, T., Blancard, B.R.S., Rozet, F., Diaconu, C., Meyer, L., Wong, K.W., Sotoudeh, H., Bietti, A. and Espejo, I. (2025) 'Walrus: a cross-domain foundation model for continuum dynamics', arXiv, arXiv:2511.15684. Available at: https://arxiv.org/abs/2511.15684 (Accessed: 2 November 2026).

Parker, L., Lanusse, F., Shen, J., Liu, O., Hehir, T., Sarra, L., Meyer, L., Bowles, M., Wagner-Carena, S., Qu, H. and Golkar, S. (2026) 'AION-1: omnimodal foundation model for astronomical sciences', Advances in Neural Information Processing Systems, 38, pp. 95386-95428.

UNIVERSITY OF CAMBRIDGE ai@cam

ai.cam.ac.uk

From access to adoption: case studies of AI adoption in science institutions

For questions about this report, get in touch via the information below:

Email:
contact@ai.cam.ac.uk

Visit:
ai.cam.ac.uk

Connect on X:
@ai_cam_mission

LinkedIn:
@ai-cambridge

In partnership with:

Bennett School of Public Policy Cambridge