AI for Science Strategy

Foreword by Minister for AI and Online Safety and Minister for Science, Innovation, Research and Nuclear

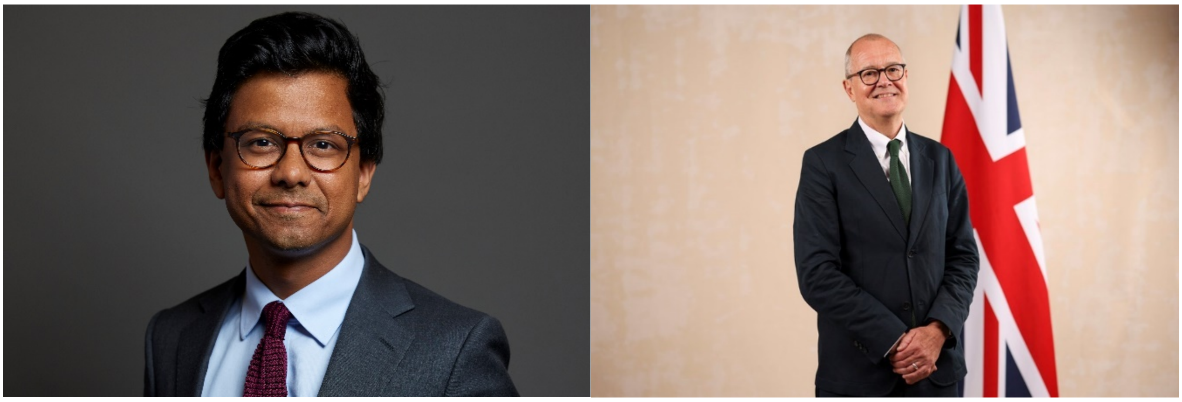

Nothing reflects the UK's capacity for innovation and creativity more strongly than our centuries-long history of transformative scientific discovery. The laws of motion, natural selection, vaccination, antibiotics, the structure of DNA – we have made discoveries that changed the world, making us healthier, safer, happier, and more prosperous. The UK's role in advancing artificial intelligence (AI) has continued this heritage. Turing, Babbage and Lovelace provided the foundations on which modern AI stands, and modern pioneers like Demis Hassabis and John Jumper have demonstrated its capacity to accelerate scientific discovery in areas like protein folding. These breakthroughs aren't just about individuals, they're also about institutions – the universities, research institutes, academies and private sector companies that have nurtured scientific excellence for generations and will continue to do so in this new era.

The worlds of artificial intelligence and scientific research are already converging, with enormous implications for the future of science. For the UK, that presents both a huge opportunity and a profound risk. If we move fast, we have a chance to supercharge our scientific productivity and establish UK leadership during a period of unprecedented scientific innovation. Too slow, and without a unified approached, and our scientific institutions could fall behind – leapfrogged by more ambitious and agile emerging leaders. Beyond research productivity, the potential for new companies and rapid growth in areas from pharmaceuticals to material science is both real and imminent.

With this in mind, our AI for science vision is centred around: developing a data landscape that facilitates transformative research; ensuring the right researchers have access to compute resource at sufficient scale; building research communities made up of truly interdisciplinary teams; and ensuring we capitalise on rapid developments in autonomous laboratory infrastructure and general-purpose and specialist AI science tools. Across these priorities, the strategy sets out actions to ensure the UK's scientific ecosystem not only adapts to, but benefits from the AI for science revolution. The strategy also launches AI for science missions - bold and ambitious targets which capitalise on UK academic and industry strength and aim to supercharge scientific progress enabled by AI. The strategy kicks off by launching the first mission, focused on harnessing the technology to speed up the research of new drugs and treatments.

In January, we published the AI Opportunities Action Plan – which detailed the macro-scale foundations needed to cement UK leadership in AI and unlock economic growth. To deliver on that plan, government is investing £2 billion between 2026-2030. Our AI for Science Strategy – which will direct up to £137m of that investment – provides a powerful complement to the Action Plan, delivering on its ambitions in the context of AI and scientific discovery. The strategy complements wider interventions towards achieving our AI for Science ambitions such as our investments in compute and the Sovereign AI unit. It is also closely connected with our UK Modern Industrial Strategy, targeting five priority areas representing frontier industries and technologies across the eight industrial strategy sectors – advanced materials, nuclear fusion, medical research, engineering biology, and quantum technology.

This strategy stands as a beacon for the UK's ambitions in AI for science, signalling both the scale of our commitments and setting clear direction for the rest of the ecosystem. In taking forward the actions set out in this plan, we won't just be keeping pace with developments in AI for science – we'll be defining its future.

Kanishka Narayan MP
Minister for AI and Online Safety in the Department for Science, Innovation and Technology

Lord Vallance
Minister for Science, Innovation, Research and Nuclear in the Department for Science, Innovation and Technology (DSIT) and the Department for Energy Security and Net Zero (DESNZ)

Introduction: The AI for science opportunity

An 'AI for Science' moment is unfolding. Investment and attention are being drawn to the idea that increasing the productivity of scientific research will be the most valuable application of AI.[1] Science is being transformed at unprecedented pace.

In the last three years, the leading edge of AI for science has moved from prediction to action. AI models are increasingly autonomous participants in the scientific process. In biology, AI models have gone from predicting the structures of known proteins to designing entirely new ones, helping scientists to accelerate research in nearly every field – from fighting malaria to potential new Parkinson's treatments.[2]

Alongside increasingly capable narrow models, frontier AI labs are now racing to build AI science agents capable of automating core parts of the scientific process. These systems can generate hypotheses, design experiments and conduct analysis without direct human input. Pairings of AI with real-world experiments foreshadow systems that can 'learn from doing' in real time, as demonstrated by Liverpool's Materials Innovation Factory, which built a mobile robotic chemist that conducted 688 experiments over eight days and discovered a new catalyst without human intervention.[3] These AI-centred discovery processes could transform scientific productivity and progress.[4]

AI in science matters because supercharged science productivity is a societal shift that dramatically improves lives. AI could slash the average drug discovery timeline; improved weather forecasts and flood predictions could safeguard national infrastructure and improve food security; and AI-augmented plasma control in fusion reactors could hasten progress to limitless clean energy.

The UK ranks 4th in world for the quality of its AI research[5] and 3rd in the world as a destination for elite AI researchers to work.[6] Homegrown startups like Latent Labs, Cusp.ai, DaltonTx, and Orbital have emerged from the UK's vibrant AI for science ecosystem and could be a profound source of UK influence and economic growth. Strengthening this AI for science ecosystem is central to our strategy, and the Sovereign AI Unit will prioritise these areas for interventions that support AI companies scaling and driving growth in the UK.

Our universities and research institutions are among the strongest in the world. Across the UK they are moving to capture the AI for science opportunity to improve lives.

---

[1] OECD, Artificial Intelligence in Science: Challenges, Opportunities and the Future of Research, 2023.
[2] University of Cambridge, AI speeds up drug design for Parkinson's ten-fold, 2024; Medicines for Malaria Venture, MMV announces new partnership using AI-powered image analysis to accelerate malaria drug discovery, 2025.
[3] Burger et. al, A mobile robotic chemist, 2020.
[4] Google DeepMind, A new golden age of discovery, 2024
[5] Tortoise, The Global AI Index, 2024.
[6] MacroPolo, The Global AI Talent Tracker 2.0, 2023.

Moorfields Eye Hospital and University College London's RETFound model can not only detect sight-threatening eye illness, but also predict heart disease.[7] This example, like many others, is underpinned by developments underway at our superb research institutions, data holders, and national facilities. Organisations like the Francis Crick Institute, the Sanger Institute, the Henry Royce Institute, the Harwell Science and Innovation Campus, UK Biobank, and EMBL's European Bioinformatics Institute (EMBL-EBI) are all taking steps to transform UK research using AI.

The UK is a scientific nation. Scientific discovery is one of the core drivers of human progress. The UK must act decisively to maintain its scientific leadership and seize the opportunity to shape the transformation of science by AI. We have all the strengths and assets to do so, but we must follow the US, the EU, and others in setting an ambitious national strategy and vision that catalyses the quick action the moment requires.

[7] UCL, World-first AI foundation model for eye care to supercharge global efforts to prevent blindness, 2023.

Objectives and approach

This strategy has two objectives:

- **To develop frontier capability in AI-driven science.** The companies and researchers developing general-purpose AI science tools and building autonomous lab infrastructure are transforming the process of discovery. It is a crucial strategic area to build UK capacity.
- **To ensure the UK retains its position of global scientific leadership**. The integration of AI into science is going to reshape the national and global research landscape. We must adapt to this transformation of science to create growth and capture the benefits for public good.

The first objective will be addressed directly in the opening section on AI-driven science. The second objective will be addressed by actions across three pillars: (1) Data, (2) Compute, (3) People and culture.

The adoption of AI in science, as everywhere, will be fast in some fields and slower in others. Whilst the UK is genuinely world leading in many areas, we must accept that cannot be true in every domain. As such, the actions in this strategy will be targeted towards five broad priority areas identified on the basis of existing UK strength, alignment with wider UK strategy - including the Modern Industrial Strategy - and opportunities for AI-driven progress. They are:

(1) engineering biology
(2) fusion energy
(3) materials science
(4) medical research
(5) quantum technologies

Targeted interventions within these domains will prioritise near-term impact, creating exemplar cases to be emulated elsewhere in the ecosystem and delivering tangible impacts that are felt by citizens in the near term. Crucially, each of our priority areas are characterised by diverse methodological approaches, research communities, data landscapes, and levels of AI-integration. In practice, therefore, the delivery of interventions across will require a tailored approach that is sensitive to the needs of specific disciplines.

The strategy also launches *AI for science missions*. These are specific, timebound, ambitious targets that can only be reached through breakthrough scientific progress enabled by AI. Delivery of this programme of interventions will bring together government's critical levers – compute resource, data assets, convening power, investment and more – as well as up to £137m of targeted investment to accelerate AI-driven scientific breakthroughs in areas of UK priority.

Government is keenly aware of the importance of responsible AI adoption in science and beyond. Of particular importance is understanding how scientists will adapt to using new tools that can hallucinate or make errors, and ensuring the adoption of AI does not negatively impact research integrity. Our approach in general will be guided by the national framework on research integrity – which ensures that all UK research is underpinned by the clear set of principles, driving high quality in all research activity.[8] Specific actions will be developed with responsible AI at their core.

Managing the environmental impacts of research-related activity and reducing the carbon footprint of the UK AI infrastructure and data storage will be considered throughout this strategy. There are innovation opportunities in sustainable AI for science, and we are looking towards potential applications that increase the productivity and data richness of individual experiments and reduce dependencies on chemical and material usage to deliver an overall reduction in negative impacts across the R&D pipeline.

AI for science is driving the UK's Modern Industrial Strategy

The Digital and Technologies Sector Plan names AI as one of six frontier technologies which are essential to driving growth, and this strategy targets five priority areas representing frontier industries and technologies across the eight industrial strategy sectors – advanced materials, fusion energy, medical research, engineering biology, and quantum technology. Across government, there is recognition of the scope for AI-driven progress across these domains.

In advanced materials, this includes programmes like The National Materials Innovation Programme, with an initial £50m committed in the UK's Modern Industrial Strategy that includes funding to deliver data-driven tools that enhance materials discovery, design, and validation. AI is transforming engineering biology by enabling novel design, high throughput analysis, and reducing experimental cost. AI advancements can impact across the design, build, test, learn cycle and are accelerating the development of solutions to some of society's most urgent challenges. For instance, by unlocking new avenues for drugs that can treat diseases and sustainable biology-derived alternatives to everyday consumer items from chemicals and materials to low-carbon fuels.

In medical research, the Life Science Sector Plan makes clear that AI is revolutionising the life sciences sector across research, diagnostics, treatment, and manufacturing. Significant investment support from across UK Research and Innovation (UKRI) and National Institute for Health and Care Research (NIHR), plus the £600m announced for the development of the Health Data Research Service, will enable world class R&D in medical research. Exciting things are also happening in fusion energy and quantum

[8] Universities UK, The Concordat to Support Research Integrity, 2025.

technologies. AI will be instrumental in the UK's pursuit of commercialised fusion energy, and Culham has been announced as the UK's first AI Growth Zone. From data utilisation to materials discovery and plant operations, AI will accelerate the work of the UK's leading scientists in making fusion energy a reality. In quantum technologies, there is evidence AI can accelerate quantum system design and will in turn unlock simulation capabilities essential for future clean-energy breakthroughs.

AI-driven science

Current AI models raise the possibility of autonomous scientific reasoning. Already, AI science agents can generate novel, testable hypotheses and design experiments when given a research problem. In many disciplines the outputs must cross into the real world, and where physical experiments are needed, there is rapidly growing interest in coupling machine learning systems to robotic synthesis and characterisation in autonomous labs.

Investors and institutions drawn to this idea are moving fast. In the UK, researchers at the Whittle lab in Cambridge and Liverpool's Materials Innovation Factory are building leading AI-driven science capabilities. Companies like Lila Sciences, Future House/Edison Scientific and Periodic Labs have collectively raised almost $1bn of funding with the vision of automating the process of scientific discovery.

In some areas, achieving this extraordinary ambition seems an entirely plausible near-term possibility. The hardest steps require breakthroughs moving beyond simulated or narrowly controlled domains, real-world validation of the efficacy of AI-designed therapeutics and materials, and reproducible evidence that AI can accelerate discovery in a way that tangibly benefits scientists and citizens.

The prospect of AI automating scientific discovery is disruptive, even in limited instances, and profoundly challenges much current research practice. It is also several steps away. We should confront the implications of continued progress with a clear understanding of the capabilities and limitations of emerging AI science systems – and position the UK as a beneficiary of the change to come.

Actions

Action 1: Accelerate the development of AI-driven science in the UK.

Why it matters: AI-driven scientific progress could fundamentally alter the nature of discovery. Interfacing AI with the real world is a barrier to validation and progress in many areas and, beyond that, there is strong evidence autonomous labs will substantially accelerate scientific research in their own right.

Delivery: we will support research teams already taking rapid action to lead the way on autonomous labs, general-purpose AI science tools, and the development of AI-driven science in the UK:

- the Sovereign AI Unit will launch an open call on autonomous labs, seeking proposals to develop or scale autonomous lab platforms in the UK. This call will target the huge potential benefits to both data generation and AI-driven discovery. 'Closed loop' systems capable of analysing results in real time and using them to autonomously control further experiments will enable the generation of data that is uniquely high-quality and sensitive to the needs of model developers.

- We will work with teams building AI systems that target the full end-to-end workflow of scientific knowledge creation to develop a framework for safe deployment in the UK science ecosystem, including with the Advanced Research and Invention Agency (ARIA) through their call for exploratory 'AI scientist' proposals. This work will also consider domain specific issues of responsible and safe adoption, such as biological security implications.
- Lastly, government will explore access-models that enable safe, transparent use of general-purpose AI science tools within the UK research ecosystem.

Action 2: Fund research into the methodological implications of integrating AI into scientific research, building on the work of the UK Metascience Unit.

Why it matters: AI is not just driving new breakthroughs in science, it is fundamentally reshaping the way science is conducted, evaluated, and understood. Metascientific investigations of AI-enabled science are needed to better understand how scientific processes are changing and to ensure that methodological developments are improving, not undermining, the integrity, novelty, and quality of scientific research.

Delivery:

- UKRI and DSIT's joint UK Metascience Unit has already taken critical steps in this space to improve our understanding of how the growing adoption of AI is changing the research landscape and how governments, industry, and funding organisations should respond. The unit has funded 18 early career fellows to look at how AI is changing science, looking at topics including the use of AI in peer review, evidence synthesis, and in the generation of synthetic data, and how these changes impact the productivity, creativity and wellbeing of researchers.
- In addition to this work, the Unit will conduct a comprehensive 'National AI in Research Survey' to better understand the diffusion and adoption of AI in various aspects of the research progress across fields and career stages.

Automating discovery: The Materials Innovation Factory

The Materials Innovation Factory (MIF) is an £81 million facility co-founded by the University of Liverpool and Unilever. The MIF is the most advanced autonomous lab in Europe, using robotics and AI to conduct high-throughput, 'closed-loop' campaigns of experimentation to generate materials data at unprecedented scale. In November 2025, Liverpool announced it would build on the success of the MIF by launching an AI Materials Hub for Innovation – creating a flagship national facility in materials science.

Facilities like the MIF provide a window into the future of AI-driven science. Not only are they driving advancements in the integration of robotics and AI, they represent a profound shift in the changing human dimensions of science at the frontier. Interdisciplinary teams of chemists, roboticists, machine-learning scientists and engineers from the public and private sector work together under one roof, developing AI tools that will change the face of chemistry and materials science, as well as creating new products and ventures that fuel UK growth.

The MIF also shows that the future of AI-driven science is already growing organically across the UK, with a top British university and industry leader collaborating to file over 200 patents, create over £400m of annual sales growth,[9] and drive progress at the frontier of materials science.

Pillar 1: Data

This section views data as a key ingredient in the AI for science recipe. That means a few things: it means quality matters; it means datasets are only constitutive parts of a larger whole; and it means what's good today might not be good tomorrow. In essence, we need to take an intentional and targeted approach to generating new high-quality datasets and optimising existing data. The actions below assume our objective is not to achieve maximum data generation, but to ask: 'what data do we need to solve the challenges we face?'

On the question of whether we should generate new data or optimise existing data to make it 'AI-ready', we will prioritise the former. Newly generated datasets have the advantage of being precisely tailored to specific research challenges, and advancements in lab automation are increasing the speed at which new data can be generated. That said, there are notable exceptions when historical data cannot be replicated or is rooted in time: the decades of legacy fusion data held by UK Atomic Energy Authority (UKAEA); longitudinal data from the National Survey of Health and Development; and ocean circulation data held by the National Oceanography Centre are all examples.

High-quality datasets are the foundation of AI-enabled scientific breakthroughs and represent a growing category of internationally important strategic assets. The development of scientific datasets in areas of mutual priority is a stated objective of the recently signed UK-US Tech Prosperity Deal, and exemplifies the value of international collaboration under this pillar.

Equally, the targeted identification and hosting of high-value datasets can be a powerful locus for building UK community and stronger national ecosystems. Work here is already underway. In June, the Sovereign AI Unit provided seed funding to the OpenBind consortium to generate foundational protein-ligand structural data to power the next era of AI for drug design. Independently of the data it generates, OpenBind is a powerful

[9] UKRI, EPSRC's economic impacts from curiosity-driven research, 2025.

model for the effective consortia we need to build, and we are committed to developing that model to ensure lab automation and data generation are central to the UK's AI for science leadership.

Alongside dataset generation, storage is critical. The actions in this strategy expand upon efforts to create a streamlined model of federated data access, with data storage in proximity to our national computing centres. These developments will pay close attention to the management of sensitive data to ensure our data ecosystem meets all relevant ethical and legal obligations, as well as appropriate cybersecurity requirements.

Actions

Action 3: UKRI will establish a mandate to scale up data storage from experimental runs and simulations at UKRI-owned facilities, labs, and institutes with an aim to ensure that all relevant and useful data from every experiment is stored, curated, and made compliant with FAIR principles by 2030.

Why it matters: Alongside the need for a targeted approach to identifying high value datasets to enable scientific breakthroughs, we must also adopt an overarching 'AI first' stance to the production, storage, and optimisation of scientific data. The UK's major facilities and research organisations are already working to pioneer an AI-first approach, but leadership from government and funders is needed.

Delivery:

- DSIT and UKRI will support the UK's leading national laboratories and scientific facilities to generate AI-ready data and pioneer AI-first approaches. As a first move and signal of intent, the Science and Technology Facilities Council (STFC) will uplift the infrastructure capabilities, including those of Harwell's Diamond Light Source. This will be a critical step towards cementing 'AI first' data generation capabilities at our world class science infrastructure, and UKRI and government will explore options to extend these upgrades across the ecosystem.
- More broadly, UKRI will continue modernising its research data policy and working with scientific communities to ensure data generated through UKRI-funded research it funds aligns with FAIR (Findable, Accessible, Interoperable, and Reusable) principles, and supports a modernised, AI ready data landscape. This will include exploring the use of AI tools to support researchers in creating and optimising their data.
- UKRI is due to publish its revised data policy in 2026, and the forthcoming UKRI AI Strategy will specify details of an AI-first approach, setting out key objectives and concrete actions for normalising and standardising the development of AI-ready data.

Action 4: Identify and develop high value datasets that will unlock transformative breakthroughs in scientific priority areas.

Why does it matter? High-quality data is a key ingredient for AI-enabled scientific breakthroughs. The identification and generation of new high value datasets will be critical in unlocking the future transformative potential of AI.

Delivery:

- DSIT will develop high value datasets to unlock AI for science breakthroughs. As a starting point, DSIT will collaborate with Renaissance Philanthropy on a call to source and review datasets proposals in the scientific priority areas set out in this Strategy. We are also scaling up an initiative with the Henry Royce institute to curate and centralise high quality AI-ready materials data, building prototype repositories of standardised metadata and APIs. This will bring us a step closer to an 'AI-ready' approach to physical sciences data, which will advance the use of AI to accelerate the discovery of new materials.
- Further actions will identify high value datasets through structured discovery processes that take an actively community-led approach, and utilise external engagement, open calls, workshops and hackathons to identify the datasets our scientists need. A key priority for the programme will be the creation of new data assets for the UK.
- The data opportunities identified will partially inform the prioritisation of efforts in the overall move towards 'AI-first' data in Action 3. It will also build on our efforts in driving AI-driven science and how can we develop high-value data through automated processes.

Quality matters: The Protein Data Bank

Founded in 1971, the Protein Data Bank (PDB) is an open-access repository of protein, nucleic acid, and biomolecular structures. For decades, researchers around the globe submitted hundreds of thousands of experimentally determined structures in highly standardized formats, creating a large and extremely high-quality database. The PDB is the unsung hero of DeepMind's AlphaFold breakthrough, and it embodies the tremendous power of well-curated scientific data in the age of AI.

The original vision for the PDB was to iteratively develop a single centralised archive of known biological structures. By 2020 researchers had submitted 172,816 experimentally validated entries. DeepMind used data from the PDB to train AlphaFold, a model which predicts (rather than experimentally validates) protein structures with unprecedented accuracy. In 2021, DeepMind and EMBL-EBI published the predicted structures of 365,000 proteins. Today, the AlphaFold

Database contains predicted structures for over 200,000,000 proteins – marking an extraordinary breakthrough in structural biology and earning Demis Hassabis and John Jumper the 2024 Nobel Prize in chemistry.

AlphaFold is a landmark breakthrough in AI for science, but it remains a singular achievement. We don't have AlphaFold equivalents in materials science, fluid dynamics or neuroscience because we don't have PDB equivalents in these disciplines. The PDB was the perfect dataset fifty years in the making. By exploiting new methods of data generation, we will support UK researchers to replicate this feat – ensuring we adapt our approach towards future opportunities that AI-driven science will present.

Action 5: Launch pilot programmes for collecting 'dark data', including negative experimental data, to boost model performance in areas of UK priority.

Why it matters: The scientific data landscape is characterised by enormous quantities of 'dark data', be it unpublished experimental results, details of experimental setups, non-machine-readable data held by large institutions, or negative experimental results. Mechanisms and incentives for publishing negative data remain limited, and do not reflect the particular value of this data for machine learning. Unlocking the potential of dark data to inform future scientific discovery will be critical in developing precise, robust scientific AI models and removing positive bias.

Delivery:

- Working with large data institutions, UKRI will launch 3-5 pilot approaches that use novel methods to capture data that is currently underrepresented in large datasets – for example negative data or details of experimental setups. Pilots will be evaluated in terms of efficiency, scalability, and the usefulness of the data collected. Pilots will take a targeted approach, focusing on enhancing high value datasets identified in action 4.
- DSIT and UKRI will also engage with industry partners to seek agreements around researcher access to privately held negative data, for example in pharmaceuticals or materials.

Action 6: Build large-scale data infrastructure to host high value datasets in proximity to sovereign compute.

Why it matters: Robust, secure and accessible data storage in proximity to compute is a key resource for training AI models and will only grow in importance over the coming years. Equally, any compute system must interact seamlessly with data from across multiple, independently managed sites – often with different formats, governance, and access restrictions.

Delivery:

- Government will establish a new data repository to be co-located with Isambard-AI in Bristol and will invest in additional storage capacity at the Edinburgh Parallel Computing Centre (EPCC) – the UK's first National Supercomputing Centre and site of the UK's next national supercomputing service. Both data storage facilities will have federated capabilities, ensuring robust compute access via AI Research Resource's (AIRR) 'AIRRport' platform.
- Government will seek to continually evolve the capabilities of this federation to ensure that it provides users with the best tools to conduct transformative research and innovate. This could include the development of secure agent-to-agent interactions that manage data access within federated systems.
- Finally, our new AIRR data capability will explore how to provide secure access to high-impact health datasets through the AI Research Resource under strong privacy safeguards, enabling UK scientists to undertake more ambitious AI research using this data. It will build on the work of the Federated Research Infrastructure by Data Governance Extension (FRIDGE) programme and involve partners like UK Biobank, the world's most widely used collection of consented and de-identified human biological and health data, to define the governance and technical controls required to deploy trusted research environments (TREs) in high performance computing clusters.

Pillar 2: Compute

Compute is the engine of AI development, and effective utilisation of the UK's R&D compute infrastructure will be vital to meet the needs of our research community. Actions in this space take two forms, a) the development of hardware, software, workflows and human infrastructure needed to deliver compute access, and b) the creation of targeted allocation pathways by which that compute resource is made available to researchers at sufficient scale.

The UK is on track to develop a world-class AI compute ecosystem through the £1 billion investment in the expansion of the AI Research Resource (AIRR). Following its formal launch in July 2025, Isambard AI is one of the world's most powerful supercomputers [10] and, coupled with the Dawn supercomputer in Cambridge, represents £300m of investment. In January 2025 the Prime Minister announced the UK's first AI Growth Zone at the UKAEA's campus at Culham in Oxfordshire. Culham will house powerful AI compute suitable for fusion research and other advanced technological applications.

These investments are already bearing fruit, with UK researchers utilising sovereign compute to tackle some of the greatest research challenges of our time. This summer, the Sovereign AI Unit ran an open call for compute access at sufficient scale to support training of state-of-the-art narrow models. As a direct result, Gábor Csányi's team at Cambridge are developing their leading materials foundation model, MACE, and researchers at Imperial are developing a foundation model for health in the form of Nightingale AI, which is also supported by Engineering and Physical Sciences Research Council's (EPSRC) GenAI hub.

Moving forward, AI for science is set to be a key allocation pathway for public compute via the AIRR, which will ensure leading researchers across our priority areas have access to the compute resource they need to drive breakthroughs at the frontier.

The AIRR allocation model will include large-scale 'hero' runs for developing foundation models and small-scale allocations for PhD researchers and new adopters. We will also explore the delivery of small-scale access to AIRR through DSIT and UKRI calls which will enable researchers to experiment, test, fail, and refine models on these systems.

DSIT and UKRI will take appropriate steps to mitigate the environmental impacts of large-scale compute infrastructure. Isambard-AI, the UK's flagship AI supercomputer, is already one of the greenest supercomputers in the world. At Edinburgh Parallel Computing Centre (EPCC), the site of the next national supercomputing service, work is underway to re-use waste heat to power the heat pumps of local homes and businesses.

---

[10] Bristol, Isambard-AI is 11th fastest supercomputer in the world, 2025.

Actions

Action 7: Launch AI for Science AIRR calls that will aim to accelerate AI-enabled scientific breakthroughs in areas of UK priority.

Why it matters: Compute is a critical enabler of cutting-edge, AI-driven science, and lack of access is a common bottleneck for scientific teams operating outside of the largest companies.

Delivery: Starting in Autumn 2025, DSIT and UKRI will offer compute access via the AIRR for research in AI for science priority areas. This will build on existing access routes supporting research at all scales:

- At the smallest scale, researchers who wish to apply AI in their scientific research will be able to apply for support via the AIRR 'Gateway' route, which offers up to 10 thousand graphics processing unit (GPU) hours over three months on the Isambard-AI and Dawn supercomputers. This route is already supporting AI for science research, including a collaboration across Isambard AI, the University of Edinburgh and Imperial College London to develop a generative foundational and causal model of medical images, further supported by the EPSRC-funded Causality in Healthcare AI (CHAI) hub.
- Building on the early success of the Sovereign AI Unit's AI for science pilot call in July 2025, we will launch a broad research call inviting applications in AI for science priority areas. This will offer allocations between 200,000 – 1,000,000 GPU hours over six months to support transformative and ambitious AI-driven research in the five priority areas.
- We will also launch a call enabling 'system takeover' access for mission-focussed projects. This will offer allocation up to 1,400,000 GPU hours (~80% capacity) over two weeks to leverage our large-scale AI supercomputing capacity and to achieve scientific breakthroughs in pursuit of our missions, starting with our mission in rapid drug development.

Further iteration cycles of calls for large-scale compute access will be opened alongside future mission launches. Progress will be reviewed every 12 months to track alignment with mission objectives and effectiveness of our approach.

Utilising large-scale simulation data: Omol 25

Generating datasets from simulations through the numerical solution of the fundamental equations can enable cutting-edge AI models to be trained in fields defined by high complexity and high computational cost. This allows researchers to explore them in significantly less computationally demanding ways, transforming our understanding of the movement of galaxies, weather prediction, catalysts, design of nuclear fusion reactors, and more. [11]

Open Molecules 2025 is a high-impact recent example of such a dataset, created by Meta and the Lawrence Berkeley National Laboratory by utilising CPU downtime on high-performance computing systems. The dataset consists of over 100 million density functional theory calculations representing 83 million unique molecular systems, making it the largest of its kind by more than an order of magnitude. This rich set of chemical and structural molecular information can now be used to train more performant AI models, as exemplified by the MACE foundation models for molecules and materials recently trained on Isambard AI.

There is an opportunity for the UK to both generate and host such datasets to provide a game-changing resource to researchers in our areas of priority. With a large-scale and targeted approach, the UK could become the home of a wide range of high-value simulation datasets, creating new opportunities for international collaboration and attracting talented researchers from overseas.

Action 8: Develop a federated network of compute clusters across the UK's research-intensive ecosystem to generate high quality simulation datasets in partnership with industry and academia.

Why it matters: There is an opportunity to develop a federated network to harness compute downtime across the UK – e.g. underutilised tier 3 compute held by universities and research-intensive institutions – to generate high-value simulation datasets. Such a system would empower government to identify and generate priority datasets and convene partners to create data generation and research collaborations greater than the sum of their parts.

Delivery:

- DSIT and UKRI will work with the UK's research-intensive institutions and Jisc[12] – the UK's National Research and Education Network – to explore the feasibility of a system which enables the federated coordination of compute (including CPU, GPU, or other accelerators) downtime and capacity across the UK's national research and education network.
- Jisc's 'JANET' network provides the underlying infrastructure for this system, and there is an opportunity to develop a system which utilises that network to increase the efficiency of our compute infrastructure. DSIT and UKRI will explore project development and potential pilot programmes in the new year.

[11] Microsoft, AI4Science to empower the fifth paradigm of scientific discovery, 2022.
[12] See Jisc, 2025; GEANT, European National Research and Education Networks, 2025.

- DSIT and UKRI will convene partners across industry and academia to explore large-scale data generation and research opportunities. These datasets have the potential to function as strategically valuable assets, that drive growth and stimulate new international collaborations.

Pillar 3: People & culture

The UK's scientific research institutions are among the strongest in the world, and the depth of AI talent here is the foundation of a thriving ecosystem.

We believe the integration of AI into science demands new types of research team, organisational structures, and modes of collaboration – building on the interdisciplinary strengths that are already nurtured in UK's academic and industry ecosystem. As AI becomes an increasingly powerful driver of scientific progress, it holds the potential to create new leaders, overturn existing hierarchies, and empower new modes of science to emerge. For the UK, that represents an enormous opportunity and a profound risk.

Our goal is a rich ecosystem that:

- trains top-level scientists and technical experts,
- facilitates large-scale team-based science within research communities,
- builds deep connectivity across academia and industry to support translation and venture creation, and
- makes the UK a beacon for talented researchers all over the world.

This will strengthen the UK's ability to develop leading research and entrepreneurial talent, driving transformative scientific breakthroughs and translating these into products and services that define our future prosperity. Crucially, the development of such an ecosystem – coupled with the developments in data and compute detailed above – will create conditions which allow us to attract and retain talent in an extremely competitive international environment.

Our universities and research institutions are already responding to the cultural and organisational changes that AI in science requires – whether that's providing specialist AI training, incentivising and facilitating interdisciplinary research, supporting the career pathways of technical staff, or building strategic relationships with industry. As we embrace the AI for science opportunity, their role will continue to be critical in driving the necessary change. Universities UK – in collaboration with DSIT and UKRI – will convene senior leaders from UK universities to identify shared objectives and lessons learned on AI for science.

Actions (training and upskilling)

Action 9: Over the next five years, provision of AI for science doctoral training will expand, with an aim to initiate training of at least 1000 researchers fluent in AI and domain science or applying AI in their research.

Why it matters: To secure the UK's future leadership in AI for science it will be essential to create a pipeline for generating new researchers with expertise in both domain science and AI methods. These experts are valuable insofar as they are capable of identifying and solving AI-shaped problems in complex domains, and function as interlocutors within interdisciplinary teams.

Delivery:

- This will involve embedding AI for science in future and existing doctoral training programmes, using doctoral focal awards (DFAs) and other mechanisms to bring together universities, institutes, industry, and other partners to train PhD students.
- Emphasising interdisciplinarity and industry engagement, programmes will provide interdisciplinary training in AI and a traditional scientific domain, and directors will be encouraged to develop innovative new approaches that recognise the changing nature of advanced scientific training.

Action 10: Use interdisciplinary fellowship programmes to equip leading scientific researchers with AI capabilities and upskill communities from within.

Why it matters: UK research communities need to develop entirely new skillsets to adapt to the coming AI for science revolution, and a valuable means of driving that change is to embed highly skilled, catalytic individuals within communities through research fellowships. This also creates an opportunity to attract the world's best researchers to the UK through individual funding streams.

Delivery:

- To equip leading scientific researchers with AI capabilities and upskill communities from within. UKRI will aim to increase the number of interdisciplinary AI for science fellowships.
- Government has already funded two breakthrough fellowship schemes to drive progress in AI for science this year. The Encode AI for Science Fellowships run by Pillar VC and ARIA has embedded AI talent from around the world in research labs across the UK. Taking a different approach, the recently launched AI Pioneer Interdisciplinary Fellowship programme targets researchers without a background in core AI research, who want to build domain relevant AI capability and develop advanced AI approaches to tackle a specific research challenge in their chosen field.

Action 11: Create diverse training programmes that upskill researchers and technical experts across all disciplines and career stages.

Why it matters: As the integration of AI into science transforms the frontiers of many disciplines, a diverse AI training and upskilling ecosystem is needed to ensure early-career researchers and established academics are able to adapt to this new world. The increased use of AI means it should be become a core part of scientific training akin to statistics or the scientific method. The role of our universities and research institutions in this space cannot be overstated, and we are working with partners across universities, national institutes, funders, and industry to catalyse whole-ecosystem approach to upskilling that remains sensitive to the diverse needs of research communities.

Delivery:

- Many research organisations are already seizing the initiative on AI training: the Spärck AI scholarships have nine university partners, 38 universities have participated in government funded 'AI conversion courses', and KCL recently launched the first dedicated AI for science master's programme.
- UKRI are leading the way on training for research technical professionals and provide AI upskilling through programmes like BridgeAI, Responsible AI UK, the Science and the Hartree centre.

Actions (research structures)

Action 12: Build interdisciplinary research teams to drive AI-enabled scientific breakthroughs in areas of UK priority.

Why it matters: Above all else, AI for science research is about bringing together diverse skillsets to find innovative new methods that drive scientific progress. At scale, that means bringing together truly interdisciplinary teams of domain scientists, machine learning specialists, data wranglers, research engineers, project managers, and more.

Delivery:

- Within the headline of up to £137m AI for science allocation, we will seek to deliver long term-funding to leading researchers, with a focus on novel organisational structures that catalyse breakthroughs.
- Building on lessons from labs working at the frontier – including the AI hubs, focused research organisations, or private sector leaders like Google DeepMind and Nvidia – these investments will be characterised by large interdisciplinary teams and will seek to build critical mass investments pooling government, private sector and philanthropic funding. Over time, successful research programmes will create positive feedback loops across the ecosystem, and innovative organisational forms will become the norm.
- Building on the success of the 'hubs' model that has been developed across the National Quantum Technologies Programme Hubs and the AI Hubs, UKRI will continue to deliver grant-funding for large-scale research, supporting both interdisciplinary and cross-sectoral research.

Action 13: Invest in the full range of research technical professionals and work with the research ecosystem to develop technical career pathways.

Why it matters: Surveys of the UK ecosystem identify shortages across the full spectrum of research technical professional (RTP) roles and research institutions struggle to recruit and retain staff with vital skillsets.[13] These personnel – whether they are research software engineers, high performance computing engineers, data wranglers or data stewards – are critical to AI for science research due to the diversity of technical inputs required to develop and deploy AI tools on real-world data. Resolving these shortages will require a training pipeline of new RTPs and rewarding career pathways that limit rapid turnover of staff and effectively integrate technical staff into research teams.

Delivery:

- UKRI, working with the UK Institute for Technical Skills & Strategy (ITSS), will continue to develop new technical career pathways for research software engineers, data scientists and AI ethics specialists, ensuring they are recognised and rewarded. Building a national, interdisciplinary community of RTPs in AI will be supported through long-term funding for its Digital Research Technical Professional Skills Networks Plus, and Strategic Technical Platform investments.
- These initiatives will be complemented by investments identified in the UK Compute Roadmap, which will build a thriving technical community around the AI Research Resource.
- Universities are already collaborating to develop RTP career pathways, with the ITSS coordinating a working group of 27 universities and research institutes to address challenges RTPs face in their career development.

Action 14: Support community-driven benchmarking and evaluation to drive progress on AI for science challenges through shared assessments of model performance.

Why it matters: Formal benchmarking and evaluations establish robust and objective assessments of AI model performance in solving scientific problems. They are valuable insofar as they drive community engagement and innovation in pursuit of shared goals; prioritise standardised assessments over hype; and encourage trust in AI tools. The value of this model of community-driven evaluation is best exemplified by the Critical Assessment of Structure Prediction (CASP) benchmark in structural biology, which played a pivotal role in the AlphaFold breakthrough by charting progress in protein-folding predictions over a 25-year period.

Delivery:

- DSIT and UKRI will support the emergence of new 'CASP-style' community benchmarks for other well-suited scientific challenges. For example, progress in

[13] The Alan Turing Institute and Technopolis, Review of Digital Research Infrastructure Requirements for AI, 2023.

the development of structural materials for fusion power could be accelerated by benchmarking against reliable and chemically specific nanoscale defect data, driving accuracy gains in the predictive modelling of materials for powerplant design.

- In practical terms, this will involve convening research communities with calls to identify relevant datasets, evaluate potential methodologies, and fund the development of robust community benchmarks that drive AI for science progress. This convening work will be supported by research institutions and relevant public sector expertise.

Missions

AI for science missions are specific, timebound, ambitious targets that can only be reached through breakthrough scientific progress enabled by AI. Missions will span across our priority areas, where the UK has strengths and AI is poised to play a significant role in accelerating progress.

We are pursuing missions to:

- Ensure resources are targeted towards shared goals instead of being spread thinly across diffuse opportunities.
- Galvanise activity and coordination across UK industry and academia and attract private sector investment.
- Unlock a wide range of further scientific breakthroughs and commercialisation opportunities.

We are taking a phased approach to launching missions. Details of “Mission one” can be found below, with further missions to be selected in 2026.

Mission One

Mission One: “We will accelerate drug discovery to develop trial-ready drugs within 100 days by 2030 and contribute to deploying new treatments faster.”

Mission statement

Drug discovery is entering a transformative moment. AI is presenting opportunities to solve challenges at many stages of the drug discovery process, drastically shortening timelines to developing new treatments.[14] These opportunities range from using foundation models to search enormous spaces of drug-like chemicals to identify drug candidates, to performing ‘in silico’ clinical trials that anticipate human responses earlier and compress development timelines.[15]

The UK is positioned to be at the forefront of this transformation. Government strategies are aligned with the opportunity, including the Life Sciences Sector Plan,[16] the UK Biological Security Strategy[17] and the recently published ‘Replacing animals in science’ strategy.[18] Our data foundations in this space are strong, with established institutions like UK Biobank and EMBL-EBI and emerging assets such as OpenBind and the Health Data Research Service, reinforced by a national network of Biofoundries. We have scientific strengths in engineering biology, synthetic chemistry and health research

[14] For just one example of a recent breakthrough, see the Institute for Protein Design’s work on AI antibody design. See Bennet et. al, *Atomically accurate de novo design of antibodies with RFdiffusion*, 2025.
[15] For an overview of the AI opportunities in drug discovery, see Zhang et. al, *Artificial intelligence in drug development*, 2025.
[16] DBT, DHSC, DSIT, *Life Sciences Sector Plan,* 2025.
[17] Cabinet Office, *UK Biological Security Strategy*, 2023.
[18] DSIT, HO, DEFRA, *Replacing animals in science: a strategy to support the development, validation and uptake of alternative methods*, 2025.

underpinned by world-leading institutes such as the MRC Laboratory of Molecular Biology, the Francis Crick Institute, Manchester Institute of Biotechnology and the Wellcome Sanger Institute. And our commercial ecosystem is vibrant, with AI startups, scaleups and industrial labs operating at the frontier.

Through Mission One, our actions will focus on accelerating the preclinical phase of the drug discovery pipeline, where we believe the richest AI opportunities are concentrated. We will complement this with further actions that will aim to optimise clinical trials. By transforming the drug discovery process we will necessarily push the boundaries of AI.

This new generation of AI capabilities will create enormous opportunities for health – from tackling chronic disease, to antimicrobial resistance (AMR), to rare and new diseases – and for economic growth, making the UK a world leader in new therapeutics. While these capabilities will be readied "in peacetime", they will also be available to call on when developing responses to pathogens of pandemic potential. As such, mission success will be a significant step forward for our pandemic preparedness and will complement the goals of the '100 Days Mission'.[19]

Actions

The programme of work will be built out, but we expect Mission One to require interventions that span several stages of the preclinical drug discovery pipeline. We anticipate acting across all three pillars – data, compute, and people & culture – as well as AI-driven science to achieve the mission.

We will seek to ensure relevant UK science communities can safely deploy cutting-edge AI science agents to transform target discovery. These AI tools can pick up on signals that are scattered across vast scientific literatures and buried in omics, genetics and health datasets that no individual human can survey, and they can generate hypotheses at speed about the targets these signals imply.

We will aim to make binding affinity far more predictable, to help scientists choose and design drugs that work better with fewer off target effects. DSIT has invested £8 million in the 'OpenBind' consortium, based at Diamond Light Source, that is using breakthrough experimental technology to generate the world's largest collection of metadata-rich data on how drugs interact with proteins, enabling advancements in high quality AI models for predicting binding and early-stage drug design.

We will explore ways to improve models of developability, including immunogenicity and ADMET (absorption, distribution, metabolism, excretion and toxicity), to accelerate lead optimisation and preclinical testing. We see significant opportunities to build on work that is in train to reduce animal testing in the drug discovery process. This includes our £30m investment in a new preclinical translational research hub to bring together

[19] International Pandemic Preparedness Secretariat, *100 Days Mission Therapeutics Roadmap*, 2024.

data, cell engineering, genomic technology, and expertise to develop new in vitro and in vivo models. The hub will work closely with academia and the life sciences industry including large AI companies. We will work through the Regulatory Innovation Office (RIO), the Medicines and Healthcare products Regulatory Agency (MHRA), the Academy of Medical Sciences (AMS) and the Centres of Excellence for Regulatory Science and Innovation (CERSIs) to ensure that regulators have the capabilities to evaluate proposals in order to safely use evidence from AI-enabled ADMET models and *in silico* clinical trials within the regulatory pathway.

We will look to provide scientific teams with funding and largescale compute access via AIRR to conduct AI R&D across these opportunities, including developing more accurate models of binding prediction, immunogenicity and ADMET.

We will seek opportunities to discover next-generation technologies and programmable platforms that can be rapidly applied to new and unexpected diseases – including so-called “Disease X”.[20] These technologies may be pivotal in a future pandemic – as the mRNA technology was in the COVID-19 response.[21] This will require interdisciplinary research bringing together AI expertise with broad areas of UK scientific and biotech strength.

We will look to provide scientific teams with funding and largescale compute access via AIRR to conduct AI R&D across these opportunities. We will seek opportunities to close critical data gaps that will enable breakthrough progress in these opportunity spaces. We will explore public-private partnerships, amplifying the success of UKRI’s Prosperity Partnerships and initiatives like Open Targets.

The new Health Data Research Service (HDRS) will support these interventions and other mission-critical advancements in drug discovery. Backed by up to £600m of investment from HMG and The Wellcome Trust, HDRS will provide a single access point to large-scale health data assets from multiple sources. Linked AI-ready pathology, radiology and genomic datasets provide a foundation for AI model development and validation to identify biomarkers, facilitate disease modelling, and create pre-clinical models that improve prediction of human responses to interventions. HDRS will support clinical trial optimisation, enabling the use of AI to streamline trial design and eligibility assessments and accelerate trial delivery.

Further missions

Action 15: Select a handful of further AI for science missions to be launched in 2026.

---

[20] World Health Organization, Annual Review of the Blueprint List of Priority Diseases, 2018.
[21] An mRNA vaccine was available in record time (63 days) because the technology was developed in peacetime for other conditions – see Johns Hopkins Bloomberg School, *The Long History of mRNA Vaccines*, 2021.

DSIT and UKRI will lead on the selection of additional missions in lines with AI for Science Strategy objectives. Mission selection will leverage the Government Office for Science's capability in horizon scanning, expert engagement and science and technology analysis. Missions will be selected based on a process that integrates dialogue with academia, industry, learned academies and science funders.