HEPI

Policy Note 67 | January 2026

Using Artificial Intelligence (AI) to Advance Translational Research

Rose Stephenson, Director of Policy and Strategy at HEPI

Lan Murdock, Senior Communications Manager at Taylor & Francis

To harness the potential of AI responsibly, stakeholders must collaborate to build an ethical AI ecosystem. This involves publishers, policymakers, institutions and funders working together to establish standards and best practices.

Introduction

Translational research is a goal-oriented approach that aims to bridge the gap between scientific discovery and real-world application. It focuses on applying academic insights to solve practical problems, often requiring collaboration across disciplines. It is a critical driver of societal progress, enabling breakthroughs in medicine, technology, social sciences and the humanities to reach broader audiences and create meaningful impact.[1]

The UK produces 57 per cent more academic publications per capita than the US. However, it lags significantly behind the US in development and scale-up metrics such as business-funded research and development, patents, venture capital and unicorns (privately owned start-up companies valued at over US$1 billion).[2]

AI presents a unique opportunity to enhance translational research, offering tools to rapidly analyse large datasets, simplify complex findings into plain language summaries for audiences beyond academia and drive innovation across diverse fields. This opportunity is reflected in the ambition set out in the recent *AI for Science Strategy* to create an AI science ecosystem that supports connectivity between academia and industry, supporting the translation of scientific breakthroughs into practical applications.[3]

To consider how AI could advance the translation of discovery-based research into real-world applications, HEPI and Taylor & Francis brought together higher education leaders, researchers, AI innovators and funders to explore this topic at a roundtable event in July 2025. (The event was held under the Chatham House rule and therefore the quotes included in this paper are not attributed to individuals.)

Drawing on the roundtable discussion and case studies from UK and global research landscapes, this report uses the analogy of a river journey to examine how AI can advance translational research.

A mountain spring

In many cases of translational research – including in some of the case studies presented in this paper – the real-world problem is central at the very beginning. In other cases, this point of the research river is truly conceptual; discovery-based research. In both types of translational research, this spring can seem small and fragile. If one stands by a mountain spring, it can be difficult and sometimes impossible to see or imagine where this emerging body of water might end up.

Tributaries to the AI-supported translational research river

For the spring of an idea to gain momentum, it must be fed by strong tributaries that widen and deepen the research river. Large, accessible datasets add volume, AI expertise sharpens the flow and interdisciplinary collaboration brings new currents that merge and strengthen the direction of travel. Together, these tributaries transform a fragile beginning into a navigable channel with real potential.

Large accessible datasets

AI offers incredible power to analyse and utilise large, complex datasets to uncover patterns and insights that would take much longer to identify manually and that humans may miss.

The Nobel Prize in Chemistry 2024 recognised two AI-enabled advances in protein structure prediction: one led by David Baker, which created new proteins, and another led by Demis Hassabis and John Jumper of Google DeepMind, which predicted protein structure sequences.[4] These breakthroughs were made possible by access to extensive, open-source databases of known protein structures, principally the Protein Data Bank and large protein sequence repositories such as UniProt. The Protein Data Bank has been built over decades through mandatory data-sharing practices in the field: researchers wishing to publish newly identified protein structures must deposit the coordinates of this structure and related experimental data in an open-access protein bank. For example, the author guidelines for *Proteins: Structure, Function and Bioinformatics* require deposition at a member site of:

- the Worldwide Protein Data Bank;
- the Research Collaboratory for Structural Bioinformatics Protein Data Banks;
- the Protein Data Bank in Europe;
- the Protein Data Bank Japan; or
- the Biological Magnetic Resonance Bank.[5]

This long-standing culture of open data has created large datasets that now underpin AI-driven advances in protein science.

This dataset tributary is well-established and accessible. It is also compiled with similar methodological approaches and with comparable data, which makes them interoperable across different individual datasets.

However, this is not always the case. For some research projects, the ‘large, accessible dataset’ tributary is more of a trickle than a healthy flow of information. In some cases, the dataset remains as its own pond, unconnected to the translational research river.

Dr Gabriele Pergola, Assistant Professor in the Department of Computer Science at the University of Warwick, is experiencing this issue in his current work. The real-life problem focuses on the police response to gender-based violence and stalking. When such cases are reported, officers are required to collect and analyse digital data from the alleged victim. This is often data from their mobile phone. A mobile phone may contain five to 10 years' worth of conversations, or hundreds of thousands of messages. The stalking unit at Cheshire Police, which is piloting the use of this research on past cases, typically handles 10 cases each day. This huge pile of digital evidence is being manually analysed by the police officers, taking up endless amounts of time and resource. The team at the University of Warwick is developing an AI tool they hope can eventually be used to identify and highlight topics and conversations relevant to the investigation, including threats of harm. To train large language models you need large, accurate datasets. However, this has proven to be a stumbling block in this research. Despite recognising the importance of this work, an agreement for sharing real, anonymised conversations (between the alleged victims and perpetrators) from the police to academics developing the AI tool has not been reached. As such, the research team have been synthesising and testing their own AI-generated datasets. They have had to develop their own tributary here, rather than being able to simply channel the existing data pond into their research river.

While many academics have been using AI in their research for decades, the mainstreaming of AI is highlighting the need for access to large, complex datasets.

Dr Frances Pearl, Associate Professor in Bioinformatics at the University of Sussex, is the co-lead for data and sample reuse for Cancer Research UK's (CRUK) Data Community Support Unit (DCSU), part of Cancer Research UK's data strategy. Dr Pearl is piloting a project to develop the CRUK Data Hub – she explained that the project aims to bring together data from a large number of CRUK-funded studies – creating a cancer-specific data portal, working in collaboration with the Health Data Research Innovation Gateway.[6]

She explained that there are significant challenges for researchers in finding the data and gaining access to it and this resource aims to help solve this problem, storing the metadata for the projects with information on how to contact the Data Access Committees to enable researchers to request access to the data.

This approach enables the sharing of data, rather than making it publicly available. It also enables approaches such as federated learning (machine learning where a model is trained across multiple devices or servers holding local data, instead of gathering all that data into one central location) to be a possibility. However, to support federated learning, standardised categories must be utilised in stored datasets. Currently, data is often recorded in different ways in different settings. For example, one hospital may record the details of a 'glucose test', while another records it as a 'sugar test'. Ensuring these disparate records are stored in the same format that can then be used to harness the power of AI to drive translational research is a real challenge.

Data are the foundation upon which AI-driven research is built. The quality of data feeding AI systems directly impacts their accuracy, reliability and overall effectiveness.[7]

AI knowledge and capacity

> *Have we got the educational processes in place to support what is clearly going to be a revolutionary change?* **Roundtable attendee**

To use AI to translate research, or to speed up the process, researchers either need the knowledge and skills to develop their own AI tools or access to someone who does. In discussions about interdisciplinarity, the disciplines in increasing demand include computer science and, more specifically, AI engineering.

This was serendipitously the case at the University of Bath, where Dr Lois Player was working as a Doctoral Researcher in the Department of Psychology. Having previously worked with local government departments, Dr Player was aware these departments often collected large amounts of free-text data without the ability to analyse it. For example, one local council had collected over 5,000 free-text responses on their clean air zone initiative, but did not have the time, in-house skills or resources to analyse them.

Dr Player was working on Structural Topic Modelling (STM) – a statistical method used to group large volumes of free-text comments into distinct topics. Currently, to utilise STM effectively, you typically need both knowledge of STM and the ability to interpret qualitative data. Given this skill set is not common, especially in policy settings, Dr Player set out to automate existing STM methods so that non-specialists could analyse qualitative data with little to no expertise.

The AI tool was named DECOTA – the Deep Computational Text Analyser. The project was published in Psychological Methods, creating a buzz around the potential for this type of tool.[8] Dr Player was contacted by several government departments, think tanks and academics who were interested in utilising this resource.

As discussed in more detail in an earlier HEPI and Taylor & Francis paper, *Advancing Translational Research*, the University of Bath runs interdisciplinary Centres for Doctoral Training (CDTs).[9] This led to discussions between Dr Player and Dr Ryan Hughes, an Engineering PhD student specialising in machine learning applied to electric vehicles. Both researchers reflected on the fact that 'no one person can do this, interdisciplinarity is key – you need to bring together different experts'.

To achieve this interdisciplinary working at Bath, the CDTs run week-long 'incubators', bringing together doctoral students from multiple disciplines to work on real-world research problems. However, Dr Player and Dr Hughes were clear it was only their personal connection that meant they were discussing their projects with each other on a daily basis, which then led to the collaboration. This allowed Dr Hughes to bring his expertise to the STM project at the precise moment it was needed. They were clear there was a small and specific point on the STM research project that the AI capability tributary needed to feed into the research river. It was serendipity in this case, but it raises the question of how to ensure that researchers across disciplines can access AI expertise in a precise and timely manner.

Interdisciplinary collaboration

> *AI has an absolutely crucial role in improving accessibility, publishing plain language summaries, looking at more multimedia content, multilingual, writing policy summaries, graphical research outputs and different sorts of content types.* **Roundtable attendee**

AI offers powerful capabilities to make content more discoverable by enabling comprehensive semantic search (search functions that utilise concepts and ideas and not simply keywords, giving a more accurate result), connecting different areas of expertise and enabling better collaboration and deeper integration across diverse disciplines.

This can be realised through language and terminology translation by using real-time AI translation services to enable researchers from different linguistic backgrounds to collaborate seamlessly. Similarly, AI can make specialised knowledge more accessible by translating subject-specific and technical terms into simpler language that researchers outside their own expertise can understand.

Further, in relation to knowledge synthesis, AI can synthesise information from multiple disciplines, identifying common themes and connections.

> *We need to move from PhD programmes where people work in isolation, to actually working in teams and appreciating that other people have different areas of expertise.* **Roundtable attendee**

An example of cross-disciplinary research utilising open data and powered by AI is the EU-funded AMIGOS (Active Mobility Innovations for Green and Safe Solutions) project. This uses AI-driven real-time data to understand urban mobility challenges (such as congestion, infrastructure costs, environmental and health impacts, equity and accessibility) and co-create more sustainable and people-centred mobility solutions.[10]

Releasing the dam

When AI is used effectively, it acts like the controlled release of a dam along the translational research river. Knowledge can end up pooled behind barriers created by the sheer volume of data and the slow pace of manual analysis. AI tools can essentially open the sluice gates, allowing the data to flow into the research river rapidly, increasing the speed and reach of research.

Efficiency gains

At the Massachusetts Institute of Technology (MIT), researchers have harnessed AI to design novel antibiotics with the potential to treat drug-resistant *Neisseria gonorrhoeae*, the bacteria that causes gonorrhoea. This team used AI in two different ways:

1) Using generative AI algorithms, the team created and then screened large libraries of chemical compounds for anti-microbial properties. They discovered a fragment called F1 that appeared to have promising activity against *N. gonorrhoeae*.

2) They then generated and tested molecules containing F1 for potency against *N. gonorrhoeae*. The information below details the size of this task, made possible by utilising multiple AI tools.[11]

Generative-AI algorithms generated about seven million chemical candidates containing the F1 compounds.

↓

These were computationally screened against *N. gonorrhoeae*, resulting in 1,000 potential compounds.

↓

Researchers selected 80 of these compounds and analysed whether these could be produced by chemical synthesis vendors.

↓

Two of the compounds could be synthesised in this way and one of them, named NG1, was very effective at killing *N. gonorrhoeae* in lab tests and in mice.

Given that there are approximately five million deaths per year related to antibiotic-resistant diseases, this new approach could have a significant impact on global health.[12]

Similarly, a team at the Institute of Cancer Research (ICR) have developed an AI model called the POETIC-AI platform. This model integrates and analyses patient records, treatment information, sample results, CT and MRI scan results and genetic profiling.

The platform will allow researchers around the world to analyse large complex datasets within a trusted research environment, with the hope this will lead to the discovery of new predictive biomarkers and more individualistic cancer treatment for patients.

Dr Maggie Cheang, team leader of Integrative Genomic Analysis in Clinical Trail at the ICR, said:

> *Carrying out translational clinical research of this nature is time consuming and expensive. But what once took years can now be completed in a matter of months or even days.*[13]

Pollution in the water

Impurities introduced in the research river can contaminate everything downstream. Using AI in translational research holds great promise, but also risks polluting the water: compromised data quality, murky algorithms, ethical quandaries and biases can seep in through unnoticed side channels. Unchecked, these pollutants can cloud judgements, distort results and erode trust in the research process.

The examples of 'pollution' discussed below include:

- AI integration, data quality and academic integrity.
- Deskilling.
- Bias.
- Ethical concerns.

AI integration, data quality and academic integrity

Robust data and methodological governance have always been central to academic integrity. Including a detailed methodology in research papers also allows for scrutiny at peer review and allows repeat studies to be undertaken, safeguarding the integrity of academic outputs.

High-quality, reliable data inputs

Robust and clearly communicated methodologies

Reliable data outputs

At the roundtable dinner, academics discussed the importance of understanding the algorithms used by AI tools in order to guarantee the integrity of their research output. In some cases, including the DECOTA case study at the University of Bath, this is what led to the creation of in-house AI tools.

High-quality, reliable data inputs

Robust and clearly communicated AI algorithm

Reliable data outputs

> *The key for me is the control that we have over it. We know what the data is … It's a bit like when you write a paper, you write the introduction, you write the methods … Sometimes [with AI] it's like writing the introduction and the results of the analysis without understanding what the method was.* **Roundtable attendee**

The Royal Society report *Science in the Age of AI* discusses the challenges of reproducibility in AI-powered research. The report outlines that issues such as a lack of understanding of how AI models work and insufficient documentation of experiments are leading to a growing body of irreproducible studies. That is, independent scientists will struggle to scrutinise, replicate or reproduce experiments for future studies. At the core of the reproducibility challenge are 'opaque machine learning models' – also known as black boxes.[14]

High-quality, reliable data inputs

Opaque maching learning models (black boxes)

Irreproducible and unexplainable results

The report notes that while these challenges exist with many AI models, opacity increases when models are developed in commercial settings.[15]

In the research river, high-quality data and transparent methodologies act as natural filtration systems, keeping the water clear. However, opaque AI models can introduce pockets of murkiness that obscure how conclusions have been reached.

The Royal Society report acknowledges there is a trade-off between explainability (how clearly one can explain how the results were generated) and performance – especially involving very complex AI models.[16] This raises further questions of how the research community judges AI tools to be ethically and methodologically sound. Is a paper about a tool being approved by a peer-review process enough? As with traditional research methodologies, will some become adopted by the wider community and then fall out of fashion as something better comes along? What processes and platforms are in place to make 'approved' tools available to other research teams?

There is a role for funding bodies and publishers here to outline clear expectations for how AI has been used in research methods and how this should be clearly stated in research publications.

Deskilling

Maintaining strong human judgement and critical thinking skills – the bedrock that defines the research river – is essential to prevent AI-enabled shortcuts from gradually undercutting core academic capabilities.

Over-reliance on AI tools may reduce critical thinking and practical skills. One recent study by Microsoft focused on knowledge workers and their perceptions of where and how AI impacts critical thinking. The simplified finding shows that users with higher confidence in AI exhibit less critical thinking, while users with higher self-confidence use more critical thinking.[17]

Rose Luckin, Professor of Learner Centred Design – interviewed for this report – explains that the risk is not just reduced critical thinking in the moment, it can also be the failure to develop these skills in the first place. For early-career researchers, especially, AI can handle the 'messy middle' of research, such as data cleaning, initial analysis and literature synthesis. Using automation here may result in researchers not developing their knowledge and skills of the deep procedural processes that go on to underpin expert judgements. Professor Luckin notes:

> *Researchers cannot critically evaluate AI output if they have not learnt to do the task themselves.*

Professor Luckin argues that the AI revolution represents a pivotal moment where humans need to become more intelligent, not less. In a collection of essays on artificial intelligence, published by HEPI and the University of Southampton, Professor Luckin outlines:

> *We must harness uniquely human capabilities – metacognition, social learning and contextual adaptation – while using AI to enhance, rather than replace, human intelligence.*[18]

In conversation, Professor Luckin outlines that metacognition is the antidote to deskilling. The solution is not to avoid using AI but to pair this use with explicit metacognition training. Researchers need to be taught to ask:

- What assumptions is this tool making?
- What would I expect to see if this output were wrong?
- How would I verify this independently?

This metacognitive training turns AI from a deskilling risk into an upskilling opportunity.

Professor Luckin finishes by noting 'the expertise paradox', that is, that AI tools are most safely used by experts who can spot errors, but they are most attractive to novices who lack the skills to do tasks manually. This suggests that institutions need policies about when, in a researcher's development, AI assistance is appropriate.

Similarly, at the roundtable event, one guest cautioned that we must consider the skills we want researchers to have. Agreeing and 'protecting' this skillset, by not outsourcing it to AI, will lead to higher-quality research and greater impact.

Bias

The paper 'Ethical and Bias Considerations in Artificial Intelligence / Machine Learning', published in *Modern Pathology*, outlines:

> *Bias in AI refers to systematic and unfair favouritism or prejudice in AI systems, which can lead to discriminatory outcomes. Three broad factors are responsible for biases in AI models: (1) data bias, which is the use of unrepresentative data; (2) development bias, which is the result of the inappropriate use of AI algorithms in model development; and (3) interaction bias, which is the result of improper user interactions with the model.*[19]

Datasets used to train AI models may inadvertently encode societal biases, reflecting historical inequalities or systemic injustices inherent to the data-collection process.[20] These pollutants can drift unnoticed until they accumulate, altering the river's chemistry and disadvantaging certain communities downstream.

The *Transforming Collections* project led by the University of the Arts, London, in partnership with the Tate and funded by the Arts and Humanities Research Council, aimed to tackle this issue. Historical metadata (data about the artworks) has, in many institutions, including at the Tate, reflected historical biases in who or what is described and by whom. There can also be issues of bias relating to what is not described in the artworks, including racial biases.[21]

A machine learning tool known as the Collections Transformer was developed.

The final report of this project outlines:

> *Transforming Collections sought to surface suppressed histories, amplify marginalised voices, and re-evaluate artists and artworks long ignored or side-lined by dominant narratives and institutional practices. The interdisciplinary approach brought together academic and artistic research into collections and museum practices, combined with participatory interactive machine learning (ML) design. The ML development was shaped and driven by researchers' case studies and questions, [in]terrogating small, bespoke, 'messy' datasets as well as larger collections' data. The focus was not on achieving a technical solution to address problems in collections, but on developing lightweight adaptable tools that can support their critical analyses. The resultant critical analytical tool (ML CAT or Collections Transformer) has the potential to aid the rethinking of habitual formulations, hierarchies and values expressed in collections' text-based digital records by offering critical prompts; while the creation of dynamic categorisations or tags refined by the user (that would not otherwise be made visible through standard search functions within collections databases), can surface unexpected connections and relations.*[22]

The *Transforming Collections* project resulted in a major public programme, *Museum x Machine x Me*, across Tate Modern and Tate Britain, where project insights and findings were shared with wide-ranging audiences.[23]

While bias can exist within the data used to train AI models (and therefore can be inherent in the outputs), AI is also being used to tackle bias, as demonstrated in this example of a 'curator in the loop' approach to AI-supported translational research.

Again, transparency is key here. As with the above example and examples of medical AI tools that are developed with data from only a select demographic pool of patients, there are good reasons for being aware of and accounting for historical inaccuracies and gaps in data. However, when looking at an old issue with a fresh perspective to address a previous bias, there is a risk of simply applying a different or more modern form of bias in these models.

Ethical concerns

Ensuring the ethical use of AI is vital, particularly in sensitive fields like medicine. The Council of Europe has identified six types of ethical challenges in AI:

1) **Inconclusive evidence:** Algorithms produce probable yet uncertain knowledge.
2) **Inscrutable evidence:** When generating knowledge, it is reasonable to expect the connection between input data and the conclusion should be open to scrutiny. Given the complexity and scale of some AI systems, this is not always the case.
3) **Misguided evidence:** AI-generated conclusions will only be as reliable as the data fed into the tool.
4) **Unfair outcomes:** This includes outcomes that may be inherently biased.
5) **Transformative effects:** This relates to the subtle shifts in how the world may be conceptualised and organised in a post-AI era.
6) **Traceability:** Given the complex inputs to AI systems (human users and developers, manufacturers, deploying organisations, the systems and models themselves), it can make it difficult to detect harms and find their cause. This mix of human and technological factors raises difficult questions about how to assign responsibility and liability for the impact of AI behaviours.[24]

To harness AI's potential responsibly, stakeholders must collaborate to build an ethical AI ecosystem. This involves publishers, policymakers, institutions and funders working together to establish standards and best practices.

The UK Research Integrity Office (UKRIO) has published a set of guidelines entitled *Embracing AI with Integrity*.[25] Given the fast-paced nature of AI development, this is intended to be a 'living document' and regular reviews and updates are expected.

Debris in the river

In conversations with academics to develop this paper, it was clear that some previously relied-on frameworks are not seamlessly transitioning into the world of AI-enabled research. Even when the water in the research river is clean, these large pieces of debris can slow or divert the river. In translational research outdated intellectual property agreements and under-examined accountability norms can act like boulders lodged in the current. They do not contaminate the water itself, but they can force abrupt detours and create bottlenecks. As AI tools become more widely used, removing or redesigning these obstructions will be critical to keeping the river navigable.

Intellectual property agreements

Dr Gabriele Pergola explained why traditional Intellectual Property (IP) agreements used in translational research do not easily fit the post-AI world. While it may be clear that the university owns the IP for the tool itself, who owns the IP for the knowledge generated by the tool? What if the user group has tweaked the model? What if the user came up with a brilliant and complex set of prompts that led to a novel finding? Navigating these complexities is another hurdle when using AI in translational research.

Accountability

A roundtable attendee asked who is responsible when things go wrong. Is it the AI developer? Or the human using the AI? An example was shared of the use of AI in medicine. The paper *Clinicians risk becoming 'liability sinks' for artificial intelligence* outlines that as AI-integrated healthcare becomes more mainstream, there is a risk that clinicians will absorb legal liability for errors over which they have limited control.[26]

AI-assisted healthcare often works as follows:

Information is fed into the AI tool

The algorithm collates and analyses the information

The AI tool delivers a recommendation to the clinician

The clinician considers the recommendation alongside further information such as patient history or an examination. They can then choose to accept the AI recommendation or override it. However, they may not know, for example, that the AI model was trained using a dataset that contained less accurate information about patients from some ethnic backgrounds. Regardless, the clinician may be held responsible for the decisions they make when delivering AI-supported healthcare.[27]

The paper goes on to state:

> *Analogous to the way a 'heat sink' takes up unwanted heat from a system, the human clinician risks being used here as a 'liability sink', where they absorb liability for the consequences of the AI's recommendation whilst being disenfranchised from its decision-making process.*[28]

Arriving at the harbour and preparing for open water

Even after navigating the river – strengthened by tributaries, purified from pollutants and steered around obstacles – researchers may reach a final transition point: the harbour. Here, the river opens into the wider sea of real-world application. The skills needed to traverse the research river are not always the same as those required to launch, sustain and steer a tool in the environments where it will be used. Researchers may find that, despite having navigated all the way downstream, they still need a seaworthy boat; new expertise, platforms, partnerships or commercial frameworks to carry their innovation out into the world.

Dr Hughes at the University of Bath explained the DECOTA project team were working to make their Structural Topic Modelling qualitative analysis tool accessible to non-specialists. The real-world problem they set out to solve – how colleagues in local government can analyse large datasets they already use without hiring expensive specialists – required a platform to host their analysis tool. Dr Hughes explained there were two avenues here:

1) Investing financially to develop the platform: This investment would require recuperation, leading to an increase in product prices and the tool being available only to well-resourced organisations, and therefore not the local government offices the research set out to support.

2) The tool could be made available on a competitive basis to under-resourced local governments: However, it would then prove difficult to raise the necessary investment to develop the 'front-end' of the platform, which is usable by non-specialists and can be maintained over time. It is a catch-22.

Researcher-developed AI tools have the potential to propel further translational research and to allow non-specialists to utilise academic-level techniques to undertake their own research. However, these AI-tools need to be available to the researchers (academic or otherwise) who can use them to advance their own work, and researchers need to be aware that these tools already exist.

In the academic world, as more researchers seek to use AI but need to understand the algorithms to ensure the integrity of their work, there is a risk that researchers across institutions will develop their own DECOTA-type AI models. If you replicate this across the many different AI tools that could be developed, it risks generating inefficiencies in the research community, with money being spent at different institutions developing the same or very similar models.

In the current climate of efficiency and the push towards shared services, there should be a focus on funding the open-sourcing of such tools to prevent inefficient reproduction of AI research tools across the system. Institutions and funders need to ensure funding pools are available to host such tools. Further, consideration should be given to providing access to these tools by colleagues outside of the academic research field, such as local government data analysts, to ensure that specialist analysis can be at the fingertips of more users.

Conclusion

AI has the potential to transform the translational research landscape, deepening the river of discovery and accelerating its flow toward real-world impact. Realising these benefits, however, will depend on responsible implementation, clean and accessible data, transparent methods and the removal of structural obstacles that slow the current. By investing in interdisciplinary expertise, ethical governance and the infrastructure needed to support and share AI-enabled tools, the UK can strengthen the entire research ecosystem. With the right commitments, more ideas will be able to travel the full course from concept to real-world application, delivering meaningful benefits for society.

Recommendations

Advancing ethical AI development and implementation requires collaboration across disciplines, industries and regions. The following recommendations are written for a UK-based audience; however, they may also be relevant to other global contexts.

Research Funding bodies

Funders are uniquely positioned to shape the ethical AI ecosystem by leveraging their influence, resources and networks.

Funders should:

- Consider requiring adherence to the UK Research Integrity Office (UKRIO) guidance as a condition of funding for research involved in AI.

- Develop targeted policies and legal frameworks to guide the responsible use of AI in translational research. These frameworks should promote transparency and address challenges in data quality, integrity and reproducibility.
- Prioritise investments in projects that focus on ethical AI, such as:
  - supporting research to reduce algorithmic bias and improve fairness;
  - funding efforts to make AI systems more interpretable, transparent and reproducible;
  - funding for trustworthy AI which includes a range of ethical AI solutions such as AI procurement guidelines, AI impact assessments and AI audit frameworks;[29] and
  - allocating grants specifically for interdisciplinary research that combines AI with ethics, sociology and law.
- Provide funding for platforms or databases of researcher-developed, transparent AI research models that support further translational research.
- Through the Research Excellence Framework, provide incentives for the use of open-source datasets or data deposited in such repositories.

Institutions

Institutions should:

- Implement mandatory training for research colleagues on the UKRIO *Embracing AI with Integrity* guidance. This should be tailored to different roles, including researchers and ethics reviewers.
- Embed the UKRIO *Embracing AI with Integrity* guidance into ethics review processes, research integrity frameworks and AI governance processes.
- Better reward interdisciplinary teamwork and recognise the differing skills, outputs, activity types and different contributors to such work.
- Consider incorporating translational research practices into their criteria for hiring and promotion.

Publishers

Publishers should:

- Develop and keep up to date AI policies that set clear expectations for authors, editors and peer reviewers on how to use and attribute AI in their work, for example, the Taylor & Francis AI Policy.[30]
- Guided by human oversight and expertise, promote interdisciplinary access to existing research by:
  - using AI to simplify complex findings into plain language summaries, impact assessment summaries highlighting practical applications and systematic reviews to reach broader audiences;[31]
  - using AI tools to create multimedia formats to enable different audiences beyond academia to engage with research more easily, for example, book-to-video course conversion.
- Adopt and use AI tools and standards that improve the access, visibility and impact of translational research, such as open research models, data sharing principles, data deposition support and indexing systems.[32]

- Use AI tools to increase engagement and explore research implications via AI chatbots, interactive Q&A systems and personalised learning tools.
- Under human supervision, using AI to support manuscript screening, plagiarism detection, bias detection, author accreditation, fraud detection and indexing, easing the load for editors. Publishers should be transparent about their use of AI in these processes.

Endnotes

1 Lan Murdock and Rose Stephenson, Advancing Translational Research, November 2024 **https://www.hepi.ac.uk/reports/advancing-translational-research-2/**

2 Cambridge Industrial Innovation Policy, Investment in Innovation **https://www.ciip.group.cam.ac.uk/innovation/investment-in-innovation-2024/**

3 Department for Science, Innovation and Technology, AI for Science Strategy, November 2025 **https://www.gov.uk/government/publications/ai-for-science-strategy/ai-for-science-strategy**

4 The Nobel Prize, Press Release, October 2024 **https://www.nobelprize.org/prizes/chemistry/2024/press-release/**

5 'Author Guidelines', Proteins: Structure, Function, Bioinformatics **https://onlinelibrary.wiley.com/page/journal/10970134/homepage/forauthors.html**

6 **https://healthdatagateway.org/en**

7 Data Quality in AI: Challenges, Importance & Best Practices **https://research.aimultiple.com/data-quality-ai/**

8 Lois Player et al, 'The use of large language models for qualitative research: The Deep Computational Text Analyser (DECOTA)', Psychological Methods, April 2025 **https://pubmed.ncbi.nlm.nih.gov/40193412/**

9 Lan Murdock and Rose Stephenson, Advancing Translational Research, November 2024 **https://www.hepi.ac.uk/reports/advancing-translational-research-2/**

10 Smart Cities: Integrating AI and open data **https://data.europa.eu/en/news-events/news/smart-cities-integrating-ai-and-open-data**

11 Anne Trafton, 'Using generative AI, researchers design compounds that can kill drug-resistant bacteria', MIT News, 14 August 2025 **https://news.mit.edu/2025/using-generative-ai-researchers-design-compounds-kill-drug-resistant-bacteria-0814**

12 Anne Trafton, 'Using generative AI, researchers design compounds that can kill drug-resistant bacteria', MIT News, 14 August 2025 **https://news.mit.edu/2025/using-generative-ai-researchers-design-compounds-kill-drug-resistant-bacteria-0814**

13 The Institute of Cancer Research, 'ICR Scientists unveil AI-powered platform that helps researchers speed up biomarker discoveries for personalised cancer treatment', April 2025 **https://www.icr.ac.uk/about-us/icr-news/detail/icr-scientists-unveil-ai-powered-platform-that-helps-researchers-speed-up-biomarker-discoveries-for-personalised-cancer-treatment**

14 The Royal Society, Science in the age of AI: How artificial intelligence is changing the nature and method of scientific research, October 2024 **http://royalsociety.org/-/media/policy/projects/science-in-the-age-of-ai/science-in-the-age-of-ai-report.pdf**

15 The Royal Society, Science in the age of AI: How artificial intelligence is changing the nature and method of scientific research, October 2024 **http://royalsociety.org/-/media/policy/projects/science-in-the-age-of-ai/science-in-the-age-of-ai-report.pdf**

16 The Royal Society, Science in the age of AI: How artificial intelligence is changing the nature and method of scientific research, October 2024 **http://royalsociety.org/-/media/policy/projects/science-in-the-age-of-ai/science-in-the-age-of-ai-report.pdf**

17 Hao-Ping Lee et al, The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers, April 2025 **https://www.microsoft.com/en-us/research/publication/the-impact-of-generative-ai-on-critical-thinking-self-reported-reductions-in-cognitive-effort-and-confidence-effects-from-a-survey-of-knowledge-workers/**

18 Rose Luckin, 'AI and human intelligence' in AI and the Future of Universities (edited by Giles Carden and Josh Freeman), HEPI Report 193, October 2025 **https://www.hepi.ac.uk/reports/right-here-right-now-new-report-on-how-ai-is-transforming-higher-education/**

19 Matthew G Hanna et al, 'Ethical and Bias Considerations in Artificial Intelligence/Machine Learning', Modern Pathology, Volume 38, Issue 3, March 2025 **https://www.sciencedirect.com/science/article/pii/S0893395224002667**

20 Matthew G Hanna et al, 'Ethical and Bias Considerations in Artificial Intelligence/Machine Learning', Modern Pathology, Volume 38, Issue 3, March 2025 **https://www.sciencedirect.com/science/article/pii/S0893395224002667**

21 Susan Pui San Lok et al, 'Final Report – Transforming Collections: Reimagining Art, Nation and Heritage', Zenodo, January 2025 **https://zenodo.org/records/14772312**

22 Susan Pui San Lok et al, 'Final Report – Transforming Collections: Reimagining Art, Nation and Heritage', Zenodo, January 2025 **https://zenodo.org/records/14772312**

23 Susan Pui San Lok et al, 'Final Report – Transforming Collections: Reimagining Art, Nation and Heritage', Zenodo, January 2025 **https://zenodo.org/records/14772312**

24 Council of Europe Portal, Common ethical challenges in AI **https://www.coe.int/en/web/human-rights-and-biomedicine/common-ethical-challenges-in-ai**

25 UK Research Integrity Office, Embracing AI with integrity: A practical guide for researchers, June 2025 **https://ukrio.org/ukrio-resources/embracing-ai-with-integrity/**

26 Tom Lawton et al, 'Clinicians risk becoming 'liability sinks' for artificial intelligence'. Future Healthcare Journal, Volume 11, Issue 1, February 2024 **https://pmc.ncbi.nlm.nih.gov/articles/PMC11025047/**

27 Tom Lawton et al, 'Clinicians risk becoming 'liability sinks' for artificial intelligence'. Future Healthcare Journal, Volume 11, Issue 1, February 2024 **https://pmc.ncbi.nlm.nih.gov/articles/PMC11025047/**

28 Tom Lawton et al, 'Clinicians risk becoming 'liability sinks' for artificial intelligence'. Future Healthcare Journal, Volume 11, Issue 1, February 2024 **https://pmc.ncbi.nlm.nih.gov/articles/PMC11025047/**

29 Allison Gardner et al, 'Ethical funding for trustworthy AI: proposals to address the responsibilities of funders to ensure that projects adhere to trustworthy AI practice', AI Ethics, Volume 2, Issue 2, June 2021, pp.277-291 **https://pmc.ncbi.nlm.nih.gov/articles/PMC8197676/**

30 Taylor & Francis, AI Policy **https://taylorandfrancis.com/our-policies/ai-policy/**

31 Taylor & Francis, How to write and publish a Plain Language Summary **https://authorservices.taylorandfrancis.com/publishing-your-research/writing-your-paper/how-to-write-a-plain-language-summary/**

32 Routledge Open Research, Open Data and Accessible Source Materials Guidelines **https://routledgeopenresearch.org/for-authors/data-guidelines**

99 Banbury Road, Oxford OX2 6JX | admin@hepi.ac.uk | www.hepi.ac.uk | @HEPI_news