

Summary output

The roles of research organisations, publishers, and funders in upholding responsible Artificial Intelligence use

UKRIO roundtable | 23 June 2026

Executive summary

Generative AI presents both opportunities and risks for peer review. While participants recognised that Artificial Intelligence (AI) may assist with defined tasks, there was strong agreement that human expertise, critical judgement and accountability must remain central to the review process. There was no consensus on whether AI should be used in peer review, with views ranging from complete prohibition to carefully governed use for limited purposes.

Participants identified confidentiality, intellectual property, copyright and personal data as immediate concerns, particularly where unpublished manuscripts or grant applications may be uploaded to external open AI tools (not an organisational-managed account).

Additional risks included inaccurate outputs, fabricated references, hidden bias, convergence of judgement, over-reliance and longer-term deskilling. AI was also seen as a potential accelerator of existing research integrity challenges.

The discussion highlighted the need for clear, proportionate and coordinated guidance, practical examples of permitted and prohibited use, transparency where AI makes a substantive contribution, secure infrastructure where AI use is authorised, and ongoing evaluation of whether AI applications genuinely improve review quality. Participants emphasised that policy alone will not be sufficient and that responsible use will depend on research culture, incentives, recognition for peer review and collaboration across the research system.

Quick guide

- [Key findings](#)
- [Recommendations from UKRIO](#)
- [Guest presentations](#)
- [Key discussion themes](#)
- [Recommendations arising from the roundtable](#)
- [List of roundtable attendees](#)
- [Further reading](#)

Key findings

Five themes emerged consistently throughout the discussion:

1. Human judgement is the defining feature of peer review

AI should remain a support tool. The reviewer must retain intellectual ownership, critical responsibility and accountability for the final assessment.

2. The boundary for AI use depends on purpose, task and risk

Participants distinguished low-risk assistance from uses that scope, generate or substitute for expert judgement. Clearer task-level boundaries are needed.

3. Confidentiality and rights cannot be treated as technical details

Unpublished material, personal information, intellectual property and copyright place significant constraints on what may be entered into AI systems.

4. AI can amplify capacity, misconduct and volume

AI may support checks that humans struggle to perform at scale, but it may also accelerate fabricated content and submission volumes.

5. Policy requires evidence, infrastructure and culture

Rules need practical guidance, research on outcomes, appropriate tools, credible accountability and a culture in which responsible use is normalised.

Recommendations from UKRIO

Drawing on the discussion, UKRIO has identified a number of priorities for supporting the responsible use of AI in peer review.

The recommendations are directed towards funders, publishers, research organisations, reviewers and researchers, recognising that responsibility for the responsible use of AI in peer review is shared across the research ecosystem.

Recommendations for each audience are provided at the end of this document. The recommendations focus on:

- preserving human expertise, independent judgement and accountability in peer review
- establishing clear and proportionate boundaries between appropriate AI support and unacceptable substitution
- protecting confidential information, personal data, intellectual property and copyright
- supporting appropriate transparency about substantive AI use
- strengthening AI literacy, practical training and peer review skills
- evaluating the effects of AI use on review quality, fairness and diversity of judgement
- encouraging greater consistency and coordination across the research sector
- addressing the workload, incentives and wider cultural factors that influence responsible practice

Guest presentations

The roundtable opened with three invited presentations offering perspectives

Presentation 1:

Use of generative AI when applying for funding

Sue Russell

Senior Funding Policy Adviser, Wellcome

Key messages:

- Wellcome distinguishes between minor AI use (no disclosure required) and substantive AI use (disclosure required).
- Generating application content with AI, nondisclosure, and inaccurate or fabricated information are unacceptable.
- Reviewers must not upload confidential applications into AI tools.
- AI may support administrative tasks but will not be used to assess applications or make funding decisions.
- Future priorities include detecting problematic AI-generated content and adapting funding assessment processes.

This presentation explored considerations for developing policies and procedures governing the use of generative AI in research funding applications. The development of Wellcome's policy for applicants and reviewers was described. The approach seeks to enable responsible innovation, support applicants who may benefit from AI and maintain trust in funding decisions, while recognising the difficulty of detecting and policing use.

The revised approach distinguishes minimal and substantive use. Translation, improvement of English, formatting and word-count assistance do not require disclosure, while more substantive uses do. Unacceptable uses include generating an entire application or sections of it, failure to disclose substantive use and established research-integrity problems such as inaccuracies, fabricated references or false claims. Where concerns cannot be examined by the funder, the host institution may be asked to investigate, and an application may be withdrawn or deferred.

For reviewers, the continuing red line is that confidential applications must not be uploaded into generative AI tools. Limited

side tasks may be permitted where the application itself is not entered, such as searches for previous research or language refinement, but outputs must be validated. Where inappropriate use has been identified, the review may be withdrawn and replaced, and applicants may be informed where exposure creates an intellectual-property risk.

Wellcome also set out limits on its own use. Generative AI would not assess application quality or support funding decisions, although secure enterprise tools might be used for administrative processing such as remit or eligibility. Future areas of consideration included detection of problematic images or data, earlier verification in the research pipeline, whether generative AI is contributing to rising application volumes, and whether funding assessment itself may need to evolve.

Presentation 2: Scholarly publishing in the age of AI

Lauren Ashby

Director of Journal Audits, Elsevier

Key messages:

- Elsevier's approach allows AI as a support tool, but authors remain responsible for their work and must disclose substantive AI use.
- Editors and reviewers must not upload unpublished manuscripts into AI tools, and scientific judgement must remain human.
- AI can support publishing workflows and integrity checks, but it informs decisions rather than making them.
- AI may increase both the benefits and risks of research publishing, including the scale of problematic submissions.
- Cross-sector collaboration, shared standards, and continuous improvement are essential for maintaining research integrity.

This presentation described how AI-related policy and integrity work sits within a broader publishing system that includes pre-publication screening, investigation and journal audit. A recurring theme was collaboration: information from post-publication work informs pre-publication signals, data-science teams support tools, and industry collaboration informs policy and training.

The presentation reported accelerating AI adoption alongside uncertainty among researchers about permissible use. Key concerns included inaccurate information, ethical issues, bias, outdated material, lack of precision and deskilling through over-reliance. Potential benefits included saving time, improving language and translation, identifying topics and improving the targeting of research questions.

Elsevier's approach permits AI as a support tool for authors, with responsibility remaining with the author and disclosure required for substantive use. The policy allows AI-supported data visualisation where the underlying data belong to the researcher but maintains restrictions on

generated research images and graphical abstracts because of fabrication and copyright risks.

For editors and reviewers, unpublished content must not be uploaded to AI tools. AI may support limited activities such as structuring or polishing a report, checking references or conducting literature searches, but scientific evaluation must remain human. Existing tools can assist with reviewer recommendations, journal-scope matching, duplicate-submission detection, technical checks, reference validation and integrity signals. These tools inform editors rather than making decisions.

A major concern was scale. AI is not only a possible source of integrity problems; it can greatly increase the volume and speed of problematic submissions. The presentation described the use of multiple signals to detect patterns associated with misconduct and networks of activity, while emphasising that post-publication correction decisions are made and verified by people. The overall message was that shared frameworks, continuous improvement and cross-sector collaboration are essential.

Resources:

[Elsevier's Generative AI policies for journals \(last updated June 2026\)](#)

Presentation 3: Considerations for the responsible use of generative AI in the peer review process

Dr Scott Summers

Associate Professor in Business Law, University of East Anglia

Key messages:

- Confidentiality, personal data, and copyright are key considerations when using AI in review processes.
- AI outputs may contain errors or bias, and thorough checking can outweigh any efficiency benefits.
- Excessive reliance on AI may undermine the value of expert peer review.
- Limited, clearly defined AI use was presented as the most practical approach.
- The research community should shape how AI is used, rather than allowing it to shape research decisions and judgement.

This presentation examined lessons from institutional work on generative AI in research and applied them to journal and funding review. It focused on the ethical and legal implications of entering material into large

language models, the purpose and value of peer review, and three broad policy options: prohibition, limited use and wide permission.

Confidential and personal information was identified as a central concern. A policy that asks reviewers not to upload such information assumes that they can identify it correctly, yet researchers may have limited understanding of what counts as personal information. This creates a need for practical guidance and training. A secure local or sector-wide tool could reduce some risks, although development, governance and cost would be substantial.

Copyright creates a related barrier. A reviewer normally does not own the manuscript or funding application and lacks permission to upload it to an AI tool. Requiring consent at submission might address part of this issue, but it would not remove the risk that unpublished information could be exposed. The presentation also addressed hallucinations, bias and the danger that reviewers may accept outputs without adequate checking. Full verification can, however, remove the efficiency that AI was intended to provide.

The presentation asked what peer review is for and why a particular

reviewer has been selected. If the value sought is expert judgement, over-reliance on AI may undermine the reason for involving that reviewer. It may frame the reviewer's attention, narrow the issues considered, weaken explanations and reduce diversity if several reviewers use systems trained on similar material.

Of the three policy options presented, outright prohibition was described as difficult to detect and enforce, while unrestricted use presented unacceptable risks. Limited, clearly scoped use was presented as the most workable current approach, supported by examples, transparency and accountability. The presentation concluded with a warning that AI could come to shape every stage of research, from application and review to analysis and publication. The research community should therefore lead its use rather than allow AI to determine the questions, judgements and language of research.

Key discussion themes

Following the guest presentations, attendees took part in an interactive group discussion focusing on the boundary between appropriate support and unacceptable substitution, the responsibilities of different organisations, transparency and accountability and the literacy and training needed for responsible practice. This discussion was guided by a set of questions prepared by UKRIO in advance of the event, informed by topics and suggestions submitted by attendees beforehand.

Preserving the purpose and value of peer review

Participants repeatedly asked what peer review is intended to provide. For many, its distinctive value lies in an expert's contextual, subjective and reasoned judgement about quality, originality, significance and contribution to a field. On this view, a fluent machine-generated report is not equivalent to peer review, even if it resembles one in form.

Some participants argued that no AI should be used in scholarly peer review because the requested contribution is specifically human expertise. Others saw a possible role for an AI-generated critique that a subject specialist could interrogate, provided it did not replace an independent human review. The difference between journal review and funding assessment was also noted, with some participants seeing more scope for carefully limited use in early-stage funding triage than in specialist assessment.

The discussion therefore highlighted a prior question for every policy: why was this reviewer selected and what value is expected from them? Permitted tools and tasks should follow from that answer rather than from a general assumption that AI use is inevitable.

“Presumably, you picked that reviewer because they are an expert in the field. What value does the generative AI review add?” Attendee quote

Support versus substitution

Participants distinguished between assistance that helps a reviewer express or test their own judgement and use that outsources the judgement itself. Examples of potentially supportive use included structuring a report, improving language, checking statistics, validating references, generating questions that challenge the reviewer's critique, or helping to assess whether data, code and reporting are sufficiently transparent.

Greater concern arose where a manuscript or application is uploaded at the start and the AI output frames the entire review, where a reviewer copies a generated report, or where AI determines novelty, quality or a funding outcome. The sequence of use was considered important. Completing an independent review before consulting a tool may preserve more critical autonomy than beginning with an AI-generated assessment.

Participants also questioned whether AI use genuinely improves efficiency. If every factual statement, citation and conclusion must be checked, the process may duplicate effort. If it is not checked, the reviewer risks submitting errors. Claims of efficiency therefore need evaluation against accuracy, added value and the time required for verification.

Confidentiality, copyright, intellectual property and data protection

There was broad concern about entering unpublished applications or manuscripts into external tools. The material may contain confidential ideas, personal information, commercially sensitive content or intellectual property. A reviewer generally does not own the material and may not have permission to reproduce or disclose it through an AI system.

Participants noted that a distinction between public and closed enterprise systems is not straightforward. A reviewer may believe that an institutional account is secure without understanding its contractual terms, retention arrangements, model training conditions or technical limits. Publishers and funders may be unable to vet the many systems available to reviewers at scale.

“We can’t vet that at scale. We can’t vet everybody’s enterprise account, and reviewers often don’t know what the terms and conditions of these tools are.” Attendee quote

Possible suggestions included vetting named tools, building secure in-platform assistants or developing shared sector infrastructure. However, even a secure tool would not resolve questions about the appropriate role of AI in judgement.

Accuracy, bias and convergence of judgement

Inaccurate summaries, fabricated or unverified references, and confident but misleading outputs were recurring concerns. Participants stressed that large language models generate patterns in language and do not understand or evaluate research in the same way as a subject expert. A reviewer may be well placed to identify mistakes within their own domain but much less able to assess generated code, methods or claims outside it.

Bias was treated as both a property of the models and a risk to the review system. Reviewers may not know what data a tool was trained on or how that affects its response. If many reviewers use similar systems, apparent diversity of assessment could narrow into repeated patterns and shared blind spots. This is particularly significant where review is intended to combine different perspectives.

The discussion also recognised that human reviewers have biases. AI should not therefore be described simply as biased in contrast with neutral humans. The concern is that machine bias may be opaque, difficult to interrogate and reproduced at scale, while also influencing the focus of the human reviewer.

Critical thinking and deskilling

Participants expressed concern that repeated reliance on AI could weaken critical thinking and reviewing skills. This risk may be especially pronounced for students and early-career researchers who have not yet developed the expertise needed to judge outputs, but it could also affect experienced researchers over time.

“The core issue is maintaining critical thinking. Deskillling is a real risk, and we are already seeing it in students. I think the next phase is that those who have used these skills in the past may experience the same problems.” Attendee quote

A related concern was the cumulative use of AI across the research lifecycle. If AI helps formulate the question, write the application, review the proposal, analyse the work, draft the publication and review the paper, research may converge around the patterns and assumptions of the systems used.

Human-centred governance should therefore consider not only individual tasks but their combined effect on research practice and knowledge production.

Participants called for attention to the longer-term consequences of adoption rather than focusing only on immediate workload pressures. Maintaining human capability should be treated as an objective of policy and training.

Scale, workload and existing weaknesses in the research system

High volumes of funding applications and manuscript submissions create real pressure, with participants observing a steady increase in both over recent years. Participants acknowledged that a strict expectation of fully manual review may be difficult to sustain if suitable reviewers cannot be recruited. AI-enabled checks may help identify issues that individuals cannot reliably detect across very large volumes.

At the same time, AI can increase the volume of low-quality or manipulated content being submitted. Participants discussed fabricated citations, paper mills, organised manipulation of review and citation processes, and incentives that reward publication quantity. AI was described as amplifying problems that pre-date generative AI systems rather than creating an entirely new integrity landscape.

This led to a broader focus on incentives and recognition. Peer review is often performed with limited formal reward and under significant time pressure. Policies on AI use will have limited effect if the system does not

address workload, recognition and the pressures that encourage shortcuts or systematic manipulation.

Transparency, disclosure and accountability

There was strong support for keeping responsibility with the human author, reviewer, editor or decision-maker. Where AI is used, the individual should remain accountable for the accuracy, integrity and reasoning of the final output.

The appropriate level of disclosure was less settled. Minimal assistance may not warrant the same reporting as substantive use, but a simple declaration may reveal little about what the tool contributed or how its output was checked. Requiring every prompt and output could be burdensome, discourage disclosure and be impossible where records have not been retained.

Participants therefore highlighted the need for proportionate disclosure that captures the tool, purpose and material contribution without creating an unmanageable audit requirement. The discussion also exposed the limits of honour systems. Declarations, researcher identifiers and institutional affiliations can all be misused, while AI-detection tools may produce false allegations. Accountability mechanisms need a defensible evidential threshold and fair procedures.

Policies that are clear, coordinated and workable

Three broad approaches were considered: prohibition, limited permission and broad permission based on the ideas presented by Dr Scott Summers. Complete prohibition communicates a clear boundary but may be difficult to detect and enforce. Broad permission carries major ethical, legal and quality risks. Much of the discussion therefore focused on limited use defined by task, risk and context.

Participants wanted guidance that explains not only principles but concrete examples. It should specify what reviewers may enter into a tool, which tasks are acceptable, what must be disclosed, how outputs should be verified, and what will happen when concerns arise. It should also recognise differences between journal and funding review and between generalist triage and subject-specialist evaluation.

“We need clear boundaries, detailed guidance and examples of what responsible use looks like in practice.” Attendee quote

Consistency across publishers, funders and research organisations was seen as desirable. Shared principles, terminology, and potentially common infrastructure could help reduce duplication, while still allowing organisations to tailor their approaches to their own context.

Evidence and evaluation

Policies and proposed solutions are evolving rapidly, yet there is still limited evidence about which interventions improve review quality, influence behaviour, reduce harm or lead to unintended consequences.

Examples discussed included trials of in-platform AI assistants, analysis of how reviewers prompt and interrogate tools, evaluation of automated reporting-guideline checks and investigation of whether AI-generated reviews add value. Participants also noted the need to assess detection methods, given the risk of false positives and the difficulty of proving inappropriate use.

Evaluation should consider review quality, accuracy, reproducibility, fairness, reviewer effort, diversity of judgement, applicant or author experience and long-term effects on skills. Organisations should be transparent about pilots and share findings where possible.

Literacy, practical training and cultural change

Participants described a gap between researchers' familiarity with AI interfaces, and their understanding of how the systems work. Important limitations include difficulties handling long documents, limited transparency about how AI systems are trained, inconsistent responses and a tendency to produce convincing answers even when they are inaccurate.

Participants highlighted the value of learning through realistic examples and practical experience.

Researchers can be asked to use a tool on a topic they know well and on one they do not, compare their ability to detect errors, or examine how they would feel if an AI-generated review were submitted about their own work. Such exercises make limitations and accountability more tangible than abstract rules alone.

“Researchers quickly identify what is wrong when an AI tool produces something in their own area of expertise. The challenge is helping them recognise that they may not be able to critique outputs in these areas they know less about. You don’t know what you don’t know.” Attendee quote

The discussion also highlighted the importance of research culture. Participants emphasised that behaviour is unlikely to change through policy alone. Responsible AI use needs to be supported by clear expectations, appropriate incentives, peer discussion, visible leadership, and education about both the capabilities of AI and the potential harms associated with its misuse.

Recommendations arising from the roundtable

The following recommendations draw together practical priorities identified during the discussion. They are offered to support further development by UKRIO and relevant stakeholders and do not assume consensus on all forms of AI use.

For funders and publishers

- Define the purpose of each stage of review and the value expected from a human reviewer before specifying acceptable AI support.
- Publish task-level examples of minimal, substantive and prohibited use, including clear rules on uploading confidential source material.
- Retain human responsibility for scientific or scholarly assessment and for funding, editorial and post-publication decisions.
- Use secure, vetted tools where AI-enabled checks are offered, and explain their function, limitations and governance to users.
- Develop proportionate declarations that capture the tool, purpose and contribution of substantive AI use without demanding unmanageable records.
- Evaluate pilots against review quality, fairness, accuracy, effort, diversity of judgement and unintended effects, and share findings where possible.
- Coordinate terminology, principles and guidance across organisations to reduce contradictory expectations.

Consider how application and submission volumes, incentives and reviewer recognition affect responsible behaviour.

For research organisations

- Provide practical training on confidentiality, personal data, intellectual property, copyright and the terms of AI services.
- Teach how large language models operate and where they can fail,

including context limitations, hallucination, opaque bias and confident presentation of error.

- Use experiential peer review learning activities that require researchers to critique outputs in familiar and unfamiliar domains.
- Build peer-review capability explicitly, especially for students and early-career researchers, so AI literacy does not replace foundational critical skills.
- Clarify the status and approved uses of institutional enterprise tools rather than assuming that a closed account makes all uses acceptable.
- Align institutional guidance with the requirements of major funders and publishers while making differences visible to researchers.

For reviewers and researchers

- Begin with the purpose of the review and ask whether AI adds a specific benefit that cannot be achieved more safely or directly.
- Do not upload unpublished manuscripts, applications or other confidential material unless the review organisation has explicitly provided an approved environment and permitted that use.
- Maintain an independent critical judgement and avoid allowing an AI-generated assessment to set the scope of the review.
- Check all generated claims, references, calculations and suggested criticisms. Do not submit material that cannot be explained and defended.
- Disclose substantive use honestly and in accordance with the relevant policy.
- Consider the cumulative effect of AI use on personal expertise, critical thinking and the diversity of scholarly judgement.

For the research sector collectively

- Undertake collaborative research on reviewer behaviour, effective safeguards, disclosure, detection, secure tools and the quality of AI-assisted review.
- Explore shared infrastructure only where it answers a clearly defined need and can meet confidentiality, rights, governance and cost requirements.
- Develop common examples and training resources that can be adapted across journal, funding and institutional contexts.
- Examine upstream incentives and systemic factors that can lead to misuse alongside AI-specific policy. Responsible peer review cannot be separated from research culture and reward systems.
- Keep policy under review as tools and practices change, while preserving stable principles of human accountability, fairness, confidentiality and scholarly judgement.

Concluding reflections

The roundtable showed that the responsible use of AI in peer review cannot be reduced to a choice between adoption and prohibition. Different review tasks carry different purposes and risks, and participants held legitimate differences about whether AI can support scholarly judgement without weakening it.

There was nevertheless substantial common ground. Confidential material should not be exposed through inappropriate tools. Humans must remain accountable for reviews and decisions. AI outputs require validation. Policies must be clear enough to use in practice. Evidence is needed before claims of quality or efficiency are accepted. Training must protect, rather than bypass, critical expertise. Finally, organisations across the research system need to work together because no single funder, publisher or institution can govern these developments alone.

The immediate challenge is to establish workable boundaries while the technology, patterns of use and evidence continue to evolve. The longer-term challenge is to ensure that AI supports a trustworthy research system

without displacing the human judgement, intellectual diversity and responsibility on which peer review depends.

UKRIO hopes that the insights and recommendations presented here will contribute to ongoing discussions across the sector and inform future policy, guidance and practice. As the use of AI in peer review continues to

evolve, ensuring that it supports rather than replaces human expertise, critical judgement and accountability will remain an important priority for the research community.

List of roundtable attendees

Name	Organisation	Role
Lauren Ashby	Elsevier	Director of Journal Audits (within Research Integrity)
Sarah Bangs	Public Library of Science (PLOS)	Senior Editor, Team Manager, Publication Ethics
Sandra Bendiscioli	European Molecular Biology Organization (EMBO)	Senior Policy Officer
Ruth Betterton	Biotechnology and Biological Sciences Research Council (BBSRC)	Senior Portfolio Manager
Muhammad Ali Chaudhry	Educate Ventures Research	Deputy Director Research
Mateja Celestina	University of the West of England, Bristol	Research Integrity Governance and Ethics Manager
Nisha Doshi	Cambridge University Press and Assessment	Director of Digital Publishing Product and AI Strategy
Sarah Dodd	Royal Academy of Engineering	Head of Grant Processes and Operations
Charlotte England	The Royal Society	Head of Grants Operations
Lydia Fulford	UK Research and Innovation (UKRI)	Senior Funding Policy Manager
Ed Gerstner	Springer Nature	Director of Research Environment Alliances
Olwen Grace	Royal Botanic Garden, Edinburgh	Deputy Director of Science (Collections)
Emma Gulseven	The Royal Society	Head of Editorial and Production
Jessica Hulce	Royal Academy of Engineering	Senior Grants Processes Manager
Robert Jackson	University of Leeds	AI Skills Developer
Emma Karoune	The Alan Turing Institute	Principal Investigator
Victoria Lee	University of Cambridge	Research Integrity and Governance Coordinator

Ruth Maisey	Nuffield Foundation	Programme Head, Education
Diana Maynard	University of Sheffield	Senior Research Fellow
Dagmar Meyer	European Research Council Executive Agency (ERCEA)	Policy Adviser
Elizabeth Moylan	Wiley	Associate Director, Research Integrity
Andrew Porter	Cancer Research UK Manchester Institute, University of Manchester	Research Integrity and Training Adviser
Sebastian Powers	UK Research and Innovation (UKRI)	Senior International Manager
James Rowland	The British Academy	Process and Peer Review Co-ordinator
Claudia Russo	The International Association of Scientific, Technical and Medical Publishers (STM)	Director, Public Affairs EU
Sue Russell	Wellcome	Senior Funding Policy Adviser
Timothée Schmude	European Centre for Algorithmic Transparency, Joint Research Centre	Scientific Trainee
Sara Schroter	British Medical Journal (BMJ)	Senior Researcher
Shelly Shochat	Karger Publishers	Research Integrity & Publication Ethics Specialist
Briony Stewart	The James Hutton Institute	Head of Business Solutions
Scott Summers	University of East Anglia	Associate Professor in Business Law
Mike Thelwall	University of Sheffield	Professor of Data Science
Samantha Trow	Coventry University	Research Ethics and Integrity Officer
Anné Verhoef	North-West University (NWU), South Africa	Director of the NWU AI Hub and Professor in Philosophy
Elena Vicario	Frontiers	Director of Research Integrity
Lorna Waddington	University of Leeds	Associate Professor and University Academic Integrity Lead
Youyou Wu	University College London	Associate Professor

Further reading

Wu, Y. (2026). No paper is perfect. Research on Research Institute. Available at: <https://researchonresearchinstitute.substack.com/p/no-paper-is-perfect>

Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S. et al. (2023). Testing of detection tools for AI-generated text. International Journal for Educational Integrity, 19, 26. Available at: <https://doi.org/10.1007/s40979-023-00146-z>

STM Association. (2026). Global reporting standard for AI disclosure in research: first consultation is open. Available at: <https://stm-assoc.org/global-reporting-standard-for-ai-disclosure-in-research-first-consultation-is-open/>

Research Professional News. (2026). Researchers already use AI. It's time to agree on how to use it responsibly. Available at: <https://www.researchprofessionalnews.com/rr-news-europe-views-of-europe-2026-5-researchers-already-use-ai-it-s-time-to-agree-on-how-to-use-it-responsibly/>

eLife. (n.d.). Reproducibility Project: Cancer Biology. Available at: <https://elifesciences.org/collections/9b1e83d1/reproducibility-project-cancer-biology>

Khalil, N.N., Reed, T.J. and Ciccozzi, M.R. (2026). ReviewBench: An Extensible Framework for Benchmarking Human and AI Manuscript Review. bioRxiv. Available at: <https://doi.org/10.64898/2026.04.17.719279>

Centre for Higher Education Transformation. (n.d.). REF-AI project. Available at: <https://chet.bristol.ac.uk/research/ref-ai-project/>

Elsevier. (n.d.). Researcher of the Future: Confidence in Research. Available at: <https://www.elsevier.com/insights/confidence-in-research/researcher-of-the-future>

UK Research Integrity Office (UKRIO). (2025). Embracing AI with Integrity: A Practical Guide for Researchers. Available at: <https://ukrio.org/wp-content/uploads/Embracing-AI-with-integrity.pdf>

Transparency Statement on AI Use

During the preparation of this summary, UKRIO used Microsoft Copilot (Version 1.25094.120.0), operating under Enterprise Data Protection, to assist in the identification and summarisation of themes from an audio transcript of the discussion. All outputs were reviewed, edited and verified by UKRIO staff, who retain full responsibility for the accuracy, interpretation and content of this publication.