# GOVERNING THE SCHOLARLY AI COMMONS





MAY 2025 Dr. Samuel A. Moore

### INTRODUCTION

In researching and writing this report I have gone on a journey across the gamut of AI agnostic, sceptic, Luddite, refusenik and cautious optimist. Such is the rapidly developing field of Artificial Intelligence research and application that it is not only difficult to keep up with this area, but also hard to understand one's own views on these technologies, their applications and their costs and benefits. As with any technology, AI can be used for good and for bad: from poorly rendered political artwork to well-curated models designed to make medical advances.

The question, then, is not whether AI is inherently good or bad but more concerning who controls it and with what motivation. If it is answerable to affected communities, ethicists and technological experts, AI may develop in a more productive way than if it is governed by the needs of shareholders and profit-seeking companies. The problem of AI governance – in the context of academic knowledge production – is the focus of this report.

In recent years, commercial publishers and information analytics companies have increased their reliance on AI-based technologies to conduct a range of tasks across the research lifecycle. From submission to publication and beyond, automated technologies are assisting with tasks relating to <u>fraud detection</u>, <u>peer review</u>, <u>article production</u>, and <u>citation analysis</u>. These technologies may be developed in house or introduced as part of the rapidly growing network of startups and companies benefitting from the huge injection of investment in this area. AI is big business and relies on grand claims about its efficacy and potential, making it especially important that affected communities are able to shape its implementation.

This report looks at a range of different strategies for governing the implementation of AI by commercial publishers across the research lifecycle – or what I'm calling the scholarly commons (as I shall explore below). Governance refers to the various forms of accountability in place to determine how these tools are developed, used and deployed at various stages of knowledge production, including the regulation, guidelines and processes that shape its overall direction. Governance may take place at the grassroots levels, or by shareholders, or through legislative frameworks, and represents a key site of struggle over technological development. This report is focused primarily on governance by research communities in the interests of academic research – research communities, including academics, librarians and university administrators, are therefore the key stakeholders and audience of this report.

Yet the governance of knowledge production is currently weighted heavily in favour of the market, which is to say that decisions to implement a technology or business model are determined by how profitable they are, how much labour they save, or how financially efficient they are. I am therefore interested in ways to keep the power of the market in check in the service of more responsible AI development.

### CONTEXT

But what do I mean by AI in knowledge production, and why am I only focusing on academic publishers and associated data analytics firms? Though it is a term without a fixed definition and is perhaps more of a political signifier than a technological one, I am using AI to refer to the computational simulation of human intelligence, enabling computers to seemingly perform tasks such as learning, reasoning, problem-solving, and decision-making.<sup>1</sup> AI can be categorized into generative and discriminative forms: Generative AI models, such as GPT-4 and DALL-E, create new text (or images, audio, etc.) by learning patterns from existing datasets and producing new outputs. In contrast, **discriminative AI** focuses on classification and decision-making, distinguishing between different inputs to determine labels or probabilities, as seen in models used for spam detection or facial recognition. Generative AI produces content, while discriminative AI analyses and categorizes it. This report focuses on both generative and discriminative forms of AI, despite the fact that much media attention is paid to the generative side of things – particularly the harms caused by relying on AI to churn out content that appears sensible and true. This kind of mass content generation has been a noteworthy feature of scientific publishing over the past few years, especially in the publication of fraudulent research that looks genuine.<sup>2</sup>

Indeed, the academic publishing industry is currently in the midst of a crisis around fraudulent research brought on largely due to their reliance on business models that prioritise a high volume of articles, reviewed as quickly as possible with recourse to as little paid human labour as possible. Publishers have reoriented their businesses towards article volume and economies of scale as their revenue becomes increasingly reliant on article processing charges for open access publishing – whereby authors pay to publish in their journals – alongside their historical subscription revenues.<sup>3</sup> A higher number of papers puts greater pressure on the whole process: more editorial checks are required at initial submission, more volunteer academics are needed to peer review these papers, and more staff is needed to handle production processes (often outsourced to low-wage labour in Global South countries). Adding to the fact that researchers need to publish increasing amounts as competition for careers intensifies, it is easy to see how academic publishers are losing sight of the papers they publish in the push for volume-based profit. When easy AI content generation is added to this mix, it becomes all the harder to trust what is being published. AI fits very well with the push for less paid human labour and so AI technologies may be introduced providing that they are 'good enough' to do the task. Marketbased assessments of what is 'good enough' will likely set the bar lower than more rigorously or democratically governed assessments.

<sup>&</sup>lt;sup>1</sup> Will Douglas Heaven, "What Is AI?," MIT Technology Review, 2024, <u>https://www.technologyreview.com/</u>2024/07/10/1094475/what-is-artificial-intelligence-ai-definitive-guide/.

<sup>&</sup>lt;sup>2</sup> Danielle Gerhard, "Detection or Deception: The Double-Edged Sword of AI in Research Misconduct," The Scientist, 2024, <u>https://www.the-scientist.com/detection-or-deception-the-double-edged-sword-of-ai-in-research-misconduct-72354</u>.

<sup>&</sup>lt;sup>3</sup> Mark A. Hanson et al., "The Strain on Scientific Publishing" (arXiv, September 27, 2023), <u>https://doi.org/10.48550/arXiv.2309.15884</u>.

In addition to the general push for volume, it is also important to note that academic publishers have moved to a more maximalist approach to business that seeks not just to make money from publications, but also from the entire research lifecycle, including research data, protocols, repository software and a host of connected elements. Academic publishing itself is controlled by a handful of large information organisations – sometimes referred to as the Big 5<sup>4</sup> – who are able to exercise large amounts of market power through their highly integrated platforms aimed at keeping users within their walled gardens for as long as possible.<sup>5</sup> This means that publishers are now much less interested in obtaining and publishing the best scientific content, when any content will do, but in attracting users to their platforms in order to amass user data for internal purposes and external monetisation. I have previously theorised this shift as a process of 'individuation through infrastructure' whereby publishers orient their technologies towards the needs of individual researchers by encouraging them to interact with their services at the expense of others. This process, as researchers such as Sarah Lamdan and Jeff Pooley explore, has significant implications for user privacy and can be theorised as a form of surveillance and control.<sup>7</sup> Researchers are continually tracked and nudged within these platforms and their data is used for predictive and evaluative purposes, both internally and externally.

There is a close relationship between publisher surveillance and AI development. Academic publishers amass data not just because they hope to use it themselves (although they do) but also to licence it further to other companies. One increasingly popular business model is to licence data for AI training, as for example Taylor & Francis has done with Microsoft.<sup>8</sup> While these deals typically take the form of content licensing, which is to say access to books and journals, they are also opaque in nature and could involve data about researcher behaviour on platforms such as their peer review feedback or the popular topics that researchers access. What this means is that data across the research lifecycle is a valuable resource for academic publishers, many have of whom have rebranded themselves as 'information analytics companies' or tech-led 'content providers' in a departure from traditional title of academic publisher. **This rebranding illustrates how all aspects of research are now open to monetisation, paving the way for AI deployment to aid this expansion.** 

<sup>&</sup>lt;sup>4</sup> Leigh-Ann Butler et al., "The Oligopoly's Shift to Open Access: How the Big Five Academic Publishers Profit from Article Processing Charges," Quantitative Science Studies 4, no. 4 (November 1, 2023): 778–99, <u>https://doi.org/10.1162/qss\_a\_00272</u>.

<sup>&</sup>lt;sup>5</sup> George Chen, Alejandro Posada, and Leslie Chan, "Vertical Integration in Academic Publishing : Implications for Knowledge Inequality," in Connecting the Knowledge Commons – From Projects to Sustainable Infrastructure : The 22nd International Conference on Electronic Publishing – Revised Selected Papers, ed. Pierre Mounier, Laboratoire d'idées (Marseille: OpenEdition Press, 2019), <u>http:// books.openedition.org/oep/9068</u>.

<sup>&</sup>lt;sup>6</sup> Samuel A. Moore, "Individuation through Infrastructure: Get Full Text Research, Data Extraction and the Academic Publishing Oligopoly," Journal of Documentation ahead-of-print, no. ahead-of-print (July 28, 2020), <u>https://doi.org/10.1108/JD-06-2020-0090</u>.

<sup>&</sup>lt;sup>7</sup> Jeff Pooley, "Surveillance Publishing," The Journal of Electronic Publishing 25, no. 1 (April 26, 2022), <u>https://doi.org/10.3998/jep.1874;</u> Sarah Lamdan, Data Cartels: The Companies That Control and Monopolize Our Information (Stanford, California: Stanford University Press, 2023).

<sup>&</sup>lt;sup>8</sup> SoA Policy Team, "The SoA Responds to Taylor & Francis Group's Sale of Data to Develop AI," The Society of Authors (blog), July 22, 2024, <u>https://societyofauthors.org/2024/07/22/the-soa-responds-to-taylor-francis-groups-sale-of-data-to-develop-ai/</u>.

## THE USE OF AI BY COMMERCIAL PUBLISHERS

Before turning to governance strategies, it is necessary to explore directly how AI is being implemented and used by publishers across the research lifecycle. I am not particularly concerned here with how efficacious its usage is, although this question is clearly important, but instead looking at where and how AI technologies are being used to shape or organise knowledge production and the subsequent risks involved. The dual pressures of the need to publish more papers while reducing reliance on paid human staff makes automation an attractive option for commercial organisations. Publishers are increasingly using AI technologies to detect research integrity issues, assist with editorial workflows, handle production tasks and offer post-publication services. Indeed, editorial systems are even being <u>coded using AI-assisted environments</u>. Others like Springer Nature hope to encourage more usage of generative AI to help in the push for 'openness'.<sup>9</sup>

At the point of submission, papers may be met with AI technologies to search for fraudulent work or ethical misconduct. For example, the publisher Wiley has launched its <u>'Papermill</u> <u>Detection Service</u>' that they claim uses AI to conduct the following tasks:

- <u>Papermill similarity detection</u> checks for known papermill hallmarks and compares content against existing papermills papers,
- <u>Problematic phrase recognition</u> flags unusual alternatives to established terms (such as tortured phrases),
- <u>Unusual publication behavior detection</u> identifies irregular publishing patterns by paper authors,
- <u>Researcher identity verification</u> helps detect potential bad actors,
- <u>Gen-AI generated content detection</u> identifies potential misuse of generative AI; and the
- Journal scope checker analyzes the article's relevance to the journal.

This service is available through Wiley's 'Research Exchange' manuscript submission system that is licensed for use by other publishers and so has the potential to become an industry standard use of AI to check for fraud. What is interesting here is how AI is itself being deployed to spot the use of AI, including so-called tortured phrases of awkward text that are the hallmark of automated translation. Wiley says that their tools are intended to assist humans with checking problematic papers although it remains to be seen if all the customers of their product will do the same.

The immediate risk from these kinds of research integrity checkers are the lack of transparency around the models used and the how the tools are deployed. It is not clear what criteria are

<sup>&</sup>lt;sup>9</sup> Niki Scaplehorn and Henning Schoenenberger, "Can AI Make Research More Open? - Impact of Social Sciences," Impact of Social Sciences - Maximizing the Impact of Academic Research (blog), April 4, 2025, <u>https://blogs.lse.ac.uk/impactofsocialsciences/2025/04/04/can-ai-make-research-more-open/</u>.

being used to determine whether a researcher is a 'bad actor', what 'irregular publishing patterns' means or even whether the parameters for these thresholds can be changed by end users. As with AI detectors more broadly, there is an embedded bias against non-native English speakers meaning that their deployment here will increase inequalities already present in the scholarly publishing through the identification of false positives. Furthermore, such detectors can be easily gamed, as researchers at Stanford's Human-Centred Artificial Intelligence Lab show, meaning the possibility of false negatives is equally high.<sup>10</sup>

Nonetheless, one would assume that because publishers like Wiley have access to thousands of articles that they own copyright to, their models should in theory be trained on more reliable research data. Again, without transparency here it is impossible to know what training data they are using to determine whether or not a paper is fraudulent. Yet Wiley itself was also embroiled in a recent scandal through its purchase of Hindawi, a publisher that has now been dissolved after huge numbers of its papers had to be retracted due to fraud. In 2023, for example, over 8000 Hindawi papers were retracted, which represented 80% of the total retractions that year.<sup>11</sup> It is not clear whether this huge corpus of retracted papers was used to train Wiley's AI detection software or whether adjustments were made after these retractions. Adding this to the fact that experts believe the extent of fraudulent publication is far greater than we understand, it seems clear that training a model on the existing literature is highly unlikely to identify fraudulent research going forward.<sup>12</sup>

Once a paper has been submitted. Al is further introduced as part of the peer review process. For example, Frontiers has developed the Artificial Intelligence Review Assistant (AIRA) to aid manuscript evaluation. As Frontiers describes: AIRA reads each paper and can currently make up to 20 recommendations in just seconds, including the assessment of language quality, the integrity of the figures, the detection of plagiarism, as well as identifying potential conflicts of interest'. Curiously, these activities are not traditionally seen as the central responsibility of peer reviewers, who are instead supposed to take a view on the quality of the research presented, including its methods, results and discussion, and make a recommendation on any changes required. Much like with Wiley's AI checker, there is little transparency around the machine learning that AIRA has undergone, nor is it possible to tell exactly which items are being checked and how. The obvious problem here is the extent to which AIRA stands in the place of the peer review process. Is it possible, for example, to use fewer human peer reviewers if the company feels that AIRA is sufficient? This push towards a mixture of human and AI-review is increasingly common and will also be the foundation of the Lifecycle Journals recently launched by Open Science Framework which requires all papers to have at least one human reviewer plus a mixture of machine-based forms of assessment.

<sup>&</sup>lt;sup>10</sup> Andrew Myers, "AI-Detectors Biased Against Non-Native English Writers | Stanford HAI," Stanford University Human-Centered Artificial Intelligence, 2023, <u>https://hai.stanford.edu/news/ai-detectors-biased-against-non-native-english-writers</u>.

<sup>&</sup>lt;sup>11</sup> Richard Van Noorden, "More than 10,000 Research Papers Were Retracted in 2023 – a New Record," Nature 624, no. 7992 (December 12, 2023): 479–81, <u>https://doi.org/10.1038/d41586-023-03974-8</u>.

<sup>&</sup>lt;sup>12</sup> Robin McKie, "The Situation Has Become Appalling': Fake Scientific Papers Push Research Credibility to Crisis Point," The Guardian, February 3, 2024, sec. Science, <u>https://www.theguardian.com/science/2024/feb/</u>03/the-situation-has-become-appalling-fake-scientific-papers-push-research-credibility-to-crisis-point.

Once an article is accepted for publication, AI is introduced to help produce articles including through typesetting, copyediting and semantic tagging, hastening the demise of skilled professions and paid human labour. Because these tasks are undertaken behind the scenes, there is less public information on how AI is used as part of the production process, although one editorial board claims that Elsevier introduced AI production (which Elsevier denies) that in part led to the editors' resignation.<sup>13</sup> While I am personally not against exploring how some production tasks can be effectively automated, the issue again is who gets to define what constitutes 'good enough' for something to be automated and how it effects employees in a company. With larger commercial publishers driven primarily by revenue generation, the bar for 'good enough' may be significantly lower than if these technologies were stewarded by research communities, especially this is already happening through the outsourcing of production tasks to volunteer academic editors, or even stopping tasks like copyediting altogether (as is now more common).

Upon publication, tagging and indexing are also now achieved with the help of Al. The online publishing platform Atypon has developed 'Atypon Autotagger' to use semantic tagging and metadata creation using 'pretrained Al'. These tags are applied upon publication and allow for better discovery within journal indexes and beyond. Similarly, the medical scientific index MEDLINE (owned by Ebsco) has been using an automated classification system for its articles since 2022. A recent article discovered that this service: 'frequently applied irrelevant or imprecise terms to publications while neglecting to apply relevant terms. As a consequence, relevant publications may be omitted from search results and irrelevant ones may be retrieved.'<sup>14</sup> This analysis suggests that commercial organisations are indeed using Al in a way that they consider 'good enough' even though almost half of the articles the authors sampled had some kind of tagging error that could impact its retrieval.

As an article is published and distributed, a host of new AI-powered search engines scrape its contents for search and discovery by their users. Companies such as <u>Consensus</u> and <u>Semantic</u> <u>Scholar</u> use AI to assist with search and produce summaries of the results with an interface akin to ChatGPT. These companies are backed by huge sums of venture capital whose investors will each look for a return on their investment through growth and eventual sale to a larger company, hastening the overall commercial stranglehold that a few companies have on this space. Larger publishing companies also produce <u>AI-generated summaries of research work</u> that they hope to monetise by having their authors pay \$49 for a summary of their article that they can use for 'marketing and communication purposes'. In addition, companies like <u>Scite</u> produce AI summaries of citation data that provide context around what a citation means – i.e., if it was confirming the original paper, disagreeing with it or simply neutral (something that is clearly inappropriate for more interpretive fields in the humanities and social sciences). Alongside these new services, existing scientific indexes have released AI 'research assistants' to help with

<sup>&</sup>lt;sup>13</sup> Kate Travis, "Elsevier Denies AI Use in Response to Evolution Journal Board Resignations," Retraction Watch (blog), January 6, 2025, <u>https://retractionwatch.com/2025/01/06/elsevier-denies-ai-use-in-response-to-evolution-journal-board-resignations/</u>.

<sup>&</sup>lt;sup>14</sup> Alexandre Amar-Zifkin et al., "Algorithmic Indexing in MEDLINE Frequently Overlooks Important Concepts and May Compromise Literature Search Results," Journal of the Medical Library Association : JMLA 113, no. 1 (n.d.): 39–48, <u>https://doi.org/10.5195/jmla.2025.1936</u>.

research queries and literature reviews, including <u>Scopus AI</u>, <u>Web of Science RA</u> and Clarivate's <u>Primo Research Assistant</u>.

Publishers are also using AI to translate previously published books for sale to new markets. Taylor & Francis, for example, writes in a press release that: 'until now, the time and resources required to translate books has meant that the majority remained accessible only to those who could read them in the original language. Books that were translated often only became available after a significant delay'.<sup>15</sup> While translation is described as being both author-led and human-reviewed, the control of these quality control processes (and whether to implement them at all) still rests with the publisher. The Society of Authors have criticised this move arguing that 'AI-generated translation is just one of the ways that AI is presenting an existential threat to creators'.<sup>16</sup> The problem here is that academics routinely sign over copyright of published material to the publisher, meaning that they lose control over these kinds of initiatives.

This summary of the use of AI across the research lifecycle is not exhaustive and instead intends to provide a flavour of how AI is being deployed in a variety of different ways by commercial publishers. These companies have bet big on AI and are hoping that the technology develops in a way that will allow them to further introduce automated technologies throughout their businesses. In each case, however, the direction of the technology used is determined unilaterally by the company in question, which is itself governed by market forces in the hope of generating revenue. Algorithms and datasets remain proprietary secrets while the strategic direction of each company is determined by revenue generation rather than scientific progress. It is therefore vital to understand whether and how these technologies may be made more accountable to research communities and broader societal actors.

<sup>&</sup>lt;sup>15</sup> "Taylor & Francis to Use AI Translation Tools to Publish Books 'Otherwise Unavailable in English," The Bookseller, accessed April 25, 2025, <u>https://www.thebookseller.com/news/taylor-francis-to-use-ai-translation-tools-to-publish-books-otherwise-unavailable-in-english</u>.

<sup>&</sup>lt;sup>16</sup> Ed Nawotka, "Taylor & Francis to Translate Books into English Using AI," <u>PublishersWeekly.com</u>, accessed April 25, 2025, <u>https://www.publishersweekly.com/pw/by-topic/industry-news/publisher-news/article/</u>97417-taylor-francis-to-translate-books-into-english-using-ai.html.

#### GOVERNANCE

Governance concerns authority and oversight: who has the ability to make decisions, under what conditions, and to whom is an organisation accountable? Or, more simply, as Nathan Snyder writes, it is the 'intersection of power and cooperation'.<sup>17</sup> Governance relates to organizational structure, rules, norms, mission, and participants, and how all these elements interact with one another. Many of these elements are shaped by forces external to the organisation in question, such as market activity and political/legislative factors, while many others are shaped by intrinsic or internal factors such as the site or resource being governed, the community that uses it and the rules it sets itself. In general, governance is a highly situated phenomenon because it emerges from the community or product that needs governing. <sup>18</sup> For example, a community vegetable garden will be governed in a vastly different way to a large non-profit organisation. The former may require less formality and a more horizontal way of working, while the latter will typically necessitate board structures, mission statements and other legal structures.

The situated nature of good governance also means that it does not scale well, so larger organisations struggle (or fail) to govern themselves adequately and are strongly shaped by regulation and shareholder needs. This is because the larger an organisation is, the more diverse and sprawling its user base may be, meaning that governance structures need to be highly thought through if they are to provide real accountability or power for an individual community. There is also a tension at the heart of governance of profit-driven companies who hope to operate with as little external oversight as necessary in the hope of fulfilling their legal obligations to return dividends to shareholders. This is particularly relevant in the context of AI which, as we have seen, is a largely commercial activity within scientific knowledge production and so meaningful governance structures may get in the way of quick or unilateral decision-making desired by companies. So while a number of frameworks and codes of practice exist to shape ethical or human-centred AI use, these are all useless if big publishers can choose to ignore them.<sup>19</sup>

In terms of the governance structures within individual academic publishers, most companies rely on standard forms of corporate governance, including boards of directors and often other stakeholders including library or technologist advisory boards. Publicly traded companies – such as Springer Nature and Elsevier – have a legal requirement to be governed by boards that are elected by shareholders, meaning that there is some form of democratic participation inherent in share ownership. Shareholders may also have the power to ensure that certain issues are presented to company boards for voting or discussion, which is one of the common ways in which corporate responsibility items can appear on corporate board agendas. For non-profit forms of governance, more rigorous approaches to accountability and reporting are required,

<sup>&</sup>lt;sup>17</sup> Nathan Schneider, Governable Spaces: Democratic Design for Online Life, 1st ed (Oakland, CA: University of California Press, 2024).

<sup>&</sup>lt;sup>18</sup> Schneider, 3.

<sup>&</sup>lt;sup>19</sup> <u>https://oecd.ai/en/ai-principles; https://www.nist.gov/itl/ai-risk-management-framework; https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai</u>

including mission statements, board structures, publicly accessible accounts and statutes and bylaws. These vary by jurisdiction and the kind of organisation in question.

Good community governance, as I am focusing on here, is rare within academic publishing and knowledge organisations. By community governance I mean actual forms of power or accountability that are distributed to different user groups associated with the creation, maintenance or use of a particular resource or product.<sup>20</sup> There is a serious lack of community governance within academic publishing and scholarly communication, not just at the corporate levels but all the way down to small organisations often staffed by 'benevolent dictators' who assume responsibility and control of an entire operation. This is despite the fact that academics provide much of the labour and content without remuneration but still do not tend to receive a stake in the overall direction of key policies or programmes. There are exceptions that include infrastructure providers like <u>ORCiD</u> or <u>Datacite</u>, but these typically tend to be organisations funded by a variety of stakeholders and so good governance is necessitated. The upshot is that good internal governance structures take time and resourcing that many organisations do not prioritise unless forced to do so. The following sections therefore look for ways in which to nurture governance over AI or to force organisations to behave ethically where possible.

<sup>&</sup>lt;sup>20</sup> I have previously written about this issue for Invest in Open Infrastructure <u>https://zenodo.org/records/</u> 7152907

## **GOVERNING AI**

Publishers are clearly aware of the need to better govern the use of generative AI technologies, especially in the context of submitting research authored by generative AI technologies. For example, Elsevier allows the use of Al-assisted technologies as part of the writing purpose to 'improve readability and language'. Springer Nature has a similar policy for authorship and also a policy reminding peer reviewers to not upload articles to generative AI tools (which would then feed them back into their training data). Other publications such as JAMA and Science allow generative AI technologies as long as their use has been declared within the manuscript. These policies are responding in part to the fact that it is very difficult to know when and how generative AI has been used in a scientific research article. Yet for the purposes of this report, what is most noteworthy about these policies is how they have been created unilaterally by the publishers themselves rather than with the broader help of other experts including librarians, researchers or even other publishers. Instead, commentators feel that publishers are looking to one another to determine their policies, as for example *Science* did by permitting the use of generative AI technologies simply because other publishers had done so (after they had previously prohibited their use).<sup>21</sup> This kind of governance is reactive and commercially minded when what is needed is proactive and determined by research communities themselves. So what are the different pathways and opportunities for good governance of AI in academic knowledge production?

#### Regulation

An instinctive response to the lack of community accountability and governance is to look to government oversight through regulation. The regulation of AI is currently a huge topic in policymaking and will continue to shape how these technologies develop in the near future and beyond. The EU in particular has positioned itself as a leader in AI regulation, seeking to intervene in the business practices of companies like OpenAI, Meta and Google by requiring more transparency through reporting requirements. A fierce diplomatic battle is underway between US lobbyists and EU politicians around the EU's AI Act that includes a Code of Practice and rules to handle 'systemic risk'.<sup>22</sup> Yet these regulations are primarily aimed at AI models such as Google's Gemini, Meta's Llama and OpenAI's GPT-4, especially their harms in misinformation and political interference, rather than more specific uses within academic knowledge production. As the bigger AI firms seek to wriggle out of regulation, there is also a risk that the AI Act will disproportionately impact open-source AI projects and approaches, particularly through onerous reporting requirements that will be much easier for large, commercial organisations. This issue has been discussed more thoroughly in a separate <u>Open Future position paper</u>.

<sup>&</sup>lt;sup>21</sup> Molly Coddington, "AI in Science Publication: The Good, the Bad and the Questionable," Drug Discovery from Technology Networks, accessed April 25, 2025, <u>http://www.technologynetworks.com/drug-discovery/</u><u>articles/ai-in-science-publication-the-good-the-bad-and-the-questionable-385650</u>.

<sup>&</sup>lt;sup>22</sup> Hannah Murphy, Barbara Moens, and Michael Acton, "Silicon Valley Fights EU Tech Rules with Backing from Donald Trump," Financial Times, February 19, 2025, sec. EU tech regulation, <u>https://www.ft.com/content/3e75c36e-d29e-40ca-b2f1-74320e6b781f</u>.

It is difficult to imagine EU policymakers prioritising the governance of AI used by academic publishing and analytic firms; instead, they are offering more generic guidelines on responsible scientific usage of generative AI, rather than direct ways to regulate it. Academic publishing is not a priority area for regulators – as has been shown frequently around publisher monopoly practices<sup>23</sup> – who are preoccupied by the bigger Silicon Valley tech firms. The publishing market itself is also more European with many of the big players operating out of the EU and UK, which makes it less likely that regulators will intervene. The UK government is less likely to regulate than the EU as Prime Minister Keir Starmer looks to align with the United States by 'turbocharging' the move to Al through investment in the private sector. At a time when oversight of AI is needed most, regulators are either caught up in bigger battles or completely uninterested in intervening in the market. The lack of interest from regulators does not mean that research communities should not lobby regulators and policymakers over AI, but they should be realistic in their aims when doing so. It is necessary, then, to look to more communityled approaches to governing AI in academic knowledge production. These may take the form of small interventions or larger scale coordinated efforts designed to place democratic power back into the hands of communities.

#### *Commons-based approaches*

The literature on AI is awash with arguments in favour of the 'digital commons', 'data commons' or 'AI commons'.<sup>24</sup> These arguments frequently allude to collective ownership of the digital infrastructures, code and/or training data associated with AI usage and uptake, although the term commons itself is frequently ill-defined. The idea of the commons refers in part to resources that are shared between particular groups of people, as in for example the common land shared by locals for grazing cattle and growing food. Yet, perhaps more importantly for my purposes, it also refers to the maintenance of these resources by groups of people (commoners) and the rules and norms for doing so. Seen in this light, the AI Commons is not just about making training data and code freely available to all who want it; instead, it concerns how communities can manage these resources including how they are used. This means that the commons is not a free-for-all public space but entails the creation of communities and well-considered democratic structures that govern long-term use.<sup>25</sup>

The distinction between the commons as a freely accessible public resource and a communitysteered collective resource is important and allows us to see how it can be useful for governing AI in academic knowledge production. If we take the view that the commons is a space of public

<sup>&</sup>lt;sup>23</sup> "Creating a Market to Replace Publisher Monopolies | Plan S," accessed April 25, 2025, <u>https://</u>www.coalition-s.org/blog/creating-a-market-to-replace-publisher-monopolies/.

<sup>&</sup>lt;sup>24</sup> Saffron Huang and Divya Siddarth, "Generative AI and the Digital Commons" (arXiv, March 20, 2023), <u>https://doi.org/10.48550/arXiv.2303.11074</u>; Alan Chan, Herbie Bradley, and Nitarshan Rajkumar, "Reclaiming the Digital Commons: A Public Data Trust for Training Data" (arXiv, May 21, 2023), <u>https://doi.org/10.48550/arXiv.2303.09001</u>; Shayne Longpre et al., "Consent in Crisis: The Rapid Decline of the AI Data Commons" (arXiv, July 24, 2024), <u>https://doi.org/10.48550/arXiv.2407.14933</u>.

<sup>&</sup>lt;sup>25</sup> See more on this in an interview with Stavros Stavrides and Massimo De Angelis: <u>https://www.e-flux.com/journal/17/67351/on-the-commons-a-public-interview-with-massimo-de-angelis-and-stavros-stavrides/</u>

access to digital outputs then any indiscriminate use can be justified in its name. For example, Meta has been found scraping copyright books from the pirate website LibGen in order to train its Llama 3 model, while libraries report being inundated by bots scraping their repository content for training materials.<sup>26</sup> In each case, the products of the web are there for the taking, legally or otherwise, because everyone should have access to our digital commonwealth (or so the argument goes). Yet this naïve view obscures the fact that not all companies are created equally, some like the publishing companies referenced here are hugely extractive profiteering multinationals. The digital commons needs to be defined robustly enough so that larger companies, who already maintain a competitive advantage, cannot simply act like lawless cowboys in the AI Wild West.

One example of this kind of approach to the digital commons is the <u>data trust</u>, which creates a democratic layer to the long-term management of datasets by mediating between individual data producer and end user. Data creators pool their rights together and allow the trustees to make decisions on their behalf about their data is used.<sup>27</sup> This approach reflects something on a shift from openness of data to commonness of data. In an open data world, companies would (and do) take data for their own use without any say on behalf of the original data creator/ subject. Yet in a world of collective data stewardship, conditions can be set on how data can be used, under what conditions or price. One could imagine a situation in which academic research communities collectively agree that their user data should be managed by a data trust and these trusts are then organised nationally or by universities. This could prevent bad actors from exploiting data while requiring best practice of the sector more broadly.

Yet the issue here is not whether solutions exist to promote good governance, but how to require that academic publishers adopt these solutions when they have little incentive to do so. This is where appeals to universities may be fruitful through their procurement processes. Academic libraries are continually negotiating agreements with the publishing industry over access to resources and payments for publishing. Alongside price, they also negotiate certain conditions around, for example, open access licensing and reporting data. The turn to AI sits well within these negotiations, which are typically managed by consortia rather than individual universities and so offer a real chance to introduce sector-defining conditions relating to control over the roll out of AI technologies. Libraries are key allies here both because publishing companies seek to compete with library services through their own product offerings (especially around research data management) and because libraries are distrustful of the business strategies of commercial publishers, particularly around open access fees. As more libraries lock horns with big publishers over issues such as rights retention and above-inflation price rises, AI could be an additional string to the library's bow in the negotiations.

One further strategy of collective governance is through the design of open scholarly infrastructures intended to compete with commercial outfits. Projects like <u>OpenAlex</u> (article indexing), <u>DSpace</u> (repository software) and <u>Zotero</u> (reference management) offer well adopted

<sup>&</sup>lt;sup>26</sup> Alex Reisner, "The Unbelievable Scale of Al's Pirated-Books Problem," The Atlantic (blog), March 20, 2025, <u>https://www.theatlantic.com/technology/archive/2025/03/libgen-meta-openai/682093/</u>.

<sup>&</sup>lt;sup>27</sup> Chan, Bradley, and Rajkumar, "Reclaiming the Digital Commons."

open-source alternatives to commercial products that position themselves as more ethical and researcher-led. Many of these projects conform to the <u>Principles of Open Scholarly Infrastructure</u> and its requirement for transparent community governance, meaning that AI policies can be designed by and for user communities and maintained for the long term, including issues related to privacy, accountability, fairness, safety and other AI harms. The existence of alternative open infrastructures for knowledge production offers a real opportunity to embed best practice around AI usage within their governance models, which with sufficient pressure from research communities could be adopted by larger commercial outfits.

#### Academic governance

But how should academics exert pressure over commercial publishers to adopt greater accountability over AI development? Where are the best entry points for this kind of governance? One of the most underappreciated forms of accountability lies with the editorial board leadership of academic publications. Even though they are often owned by publishers rather than research communities, journals themselves are steered by academic editors often with no or little remuneration. These boards are an untapped source of community power despite the fact that boards do not typically have a say in the policy decisions affecting their journals. There is nothing to stop this changing and academics forcing publishers to offer them a greater say in policies that affect their publications.

Currently, the relationships between journal editors and their publishers are strained, in part due to editorial distrust of practices like AI and automation but also due to commercial imperatives imposed by publishers requiring editors to accept more articles or use new submission systems and workflows. This has led to a number of public resignations and open letters about journal practices from editorial boards worried about the future of their journals. Retraction Watch maintains a <u>list of these resignations</u>, which have increased apace in recent years. Some editorial boards go on to start new journals with a different publisher, while others simply cease working on the journal entirely. This kind of collective power could be used to force publishers to adopt certain policies around when and where within the research lifecycle AI is used, how it is developed and who is responsible for its maintenance. They could do this as a condition of continuing to edit for the journal. If publishers fail to address these concerns, the journal can resign or stop working until the issue is resolved.

The idea behind this form of governance is that while one editorial board is unlikely to effect much change, more editorial boards doing so would exert pressure on the publishing industry to cede to researchers' demands. It would be possible for editorial boards to work together and decide in advance what their demands are, possibly by setting up a meta-organisation with representatives of a variety of editorial boards. The point here is to identify areas of collective power that already exist within academia and put them to good use. For example, much of this work could be democratically decided through learned society structures that are composed of disciplinary members and frequently comment on issues affecting their disciplines. A similar approach would be to coordinate the peer review work being undertaken for publishers. Without peer reviewers, scholarly publishing cannot operate and so a peer reviewers' union would be one way of collectively agreeing certain conditions in exchange for their labour. These could be

organised geographically or by discipline or even career stage, particularly in reference to the fact that academics receive a huge number of requests to review each year and so could be selective about where to perform this service.

More formalised structures and processes also exist to increase oversight of academic knowledge production. Academic trade unions are foundational sources of collective power and will intervene on issues that impact their membership. While currently in the UK and much of Europe unions are preoccupied by the labour conditions imposed by universities in the assault on higher education, there is no reason that academic publishers should be separate from this, especially as the assault on higher education also takes the form of AI-generated content being used to supplement or <u>replace the work of trained educators</u>. The move to AI across higher education is part of the overall devaluing of labour that takes place in the service of commercial aims. Forms of governance that realise how this logic works across multiple areas are vital for a joined-up approach.

Unions are also typically advised by groups of experts in particular issues, as for example the <u>Declaration on Research Assessment</u> (DORA) did when the University of Liverpool sought to make a number of staff redundant based on their research performance.<sup>28</sup> In this case, DORA intervened to emphasise that the bibliometric measures that Liverpool used to determine whom to make redundant were both flawed and in contravention with Liverpool's commitment to DORA. The appeal to DORA added an extra layer of legitimacy to the union's demands and so a similar organisation for governing scholarly publishing could have the same impact. Organisations such as the <u>Coalition for Advancing Research Assessment</u> and the <u>Barcelona Declaration on Open Research Information</u> would be useful allies here, each able to provide input based on their own expertise.

Ultimately, this kind of collective action would be the most productive form of governance for academic knowledge production, but it requires significant amounts of effort and coordination. It would be possible to create a democratic community-led organisation of multiple stakeholders while also pooling editorial and reviewer labour to advocate for ethical practices within the publishing industry, including but not limited to the roll out of AI technologies across the research lifecycle. This is an important area that is ripe for exploration.

#### Academic "citizen" assemblies

One final form of governance that is less combative but has potential is the citizen assembly, which is used to great effect in a range of sectors. Citizen assemblies are used to obtain input from the public on specific issues and have been used famously to break the deadlock in Ireland around legalising abortion. The academic Jack Stilgoe explores their importance in a recent article for *Science*, arguing:

In some cases, AI will be a technology that people can choose to use or not, but many AI applications may be more like unseen infrastructures (think about facial recognition,

<sup>&</sup>lt;sup>28</sup> "Responsible Research Assessment Faces the Acid Test," Nature 595, no. 7868 (July 21, 2021): 471–72, https://doi.org/10.1038/d41586-021-01991-z.

navigation algorithms, or advertising), where people are unaware that they are interacting with the technology. If they have little agency as consumers, then their role as citizens becomes more important.<sup>29</sup>

Citizen assemblies are important because they delegate real power to communities by both informing them of an issue at hand and then asking for their opinion on it, often with a process of voting. Assemblies can therefore build trust in technological development by allowing affected communities to shape its progress. They could be put to good use in the context of academic knowledge production because academics are invested in these areas despite not being experts in how publishing works. This is all the more important as, like Stilgoe argues, academics primarily interact with publishers as consumers rather than agents with real power.

One could imagine universities requiring citizen assemblies on AI developments, perhaps as part of their procurement processes, by bringing together librarians, technologists, publishers and academics from a range of disciplines. This process could have the ultimate aim of generating a binding code for AI development in academic research, something that could be used as a touchstone for AI developers and users of all kinds. Such a code would be especially beneficial in the context of the general democratic deficit in scholarly publishing and knowledge infrastructures, not just the big publishers but startups, non-profits and community-led projects too. The difficulty with this approach, however, is bringing these people together and agreeing on the rules for participation. There is a substantial body of literature on how to do this effectively, although the political conditions have to be amenable to such a process.

<sup>&</sup>lt;sup>29</sup> Jack Stilgoe, "AI Has a Democracy Problem. Citizens' Assemblies Can Help. | Science," 2024, <u>https://</u>www.science.org/doi/10.1126/science.adr6713.

## CONCLUSION

Throughout this report I have attempted to build a case for community governance over AI in a range of locations from commons-based forms of organisation such as data trusts and opensource providers, through to more activist interventions and union-based strategies. I have tried above all to seek entry points where governance may be possible across the research lifecycle, including through procurement processes and by editorial boards, all while gauging the realistic possibilities that these approaches merit success. Part of the problem with governing academic publishing is that research communities and commercial publishers are antagonistic towards one another and there is a serious power imbalance in favour of larger profit-driven companies. Academics are locked into this system even though they prop up the industry through the content they provide and editorial labour they offer.

However, we are ultimately reaching an unsustainable situation that is broader than just Al governance and stretches across to the relationship between universities and the publishing industry. Academic knowledge dissemination is controlled by a small handful of large commercial players that are not primarily answerable to academic research communities but the market at large.<sup>30</sup> In recent years, these publishers have broadened their business models from payment for subscriptions towards payment for article-processing charges and associated publishing agreements. This shift has meant that their business models are dependent on the number of articles they publish – more articles mean more profit – and they consequently publish as much as they can and as quickly as possible.<sup>31</sup> This not only exhausts volunteer editorial and reviewer labour, but also means that much more questionable research is published in the push for volume. This situation represents a fundamental breakdown in trust between publisher and academia and consequently with the research being published, a breakdown that is overall a question of governance based on the fact that academic publishing has been outsourced to a private industry when it had previously been under academic control.

Publishing, like academic research more generally, cannot work without trust. The turn to Al merely exacerbates the loss of trust between publisher and research community, meaning that a much deeper problem needs fixing. The work described here will in some senses be a stopgap, while in other senses it may push towards a greater reassessment of academia's relationship with the publishing industry. In any case, greater governance by research communities will need to happen if academic research is going to have a productive future with or without artificial intelligence. We should therefore look to fold AI governance concerns into the bigger push for community oversight over the future of academic knowledge production.

<sup>&</sup>lt;sup>30</sup> Aileen Fyfe et al., "Untangling Academic Publishing: A History of the Relationship between Commercial Interests, Academic Prestige and the Circulation of Research," 2017, <u>https://research-repository.st-andrews.ac.uk/bitstream/handle/10023/10884/Fyfe\_etal\_UntanglingAcPub\_CC.pdf?sequence=1</u>.

<sup>&</sup>lt;sup>31</sup> Hanson et al., "The Strain on Scientific Publishing."

## **ABOUT OPEN FUTURE**

<u>Open Future</u> is a European think tank that develops new approaches to an open internet that maximize societal benefits of shared data, knowledge and culture.

<u>Dr. Samuel A. Moore</u> is a Research Fellow at Open Future. He is the Scholarly Communication Specialist at Cambridge University Library and Principal Investigator of Materialising Open Research Practices in the Humanities and Social Sciences (funded by Wellcome Trust, AHRC and the Research England Development Fund). Samuel is also an Affiliated Lecturer at Cambridge Digital Humanities and a College Research Associate at King's College Cambridge.



This report is published under the terms of the <u>Creative Commons Attribution</u> <u>License</u>.