



A Third Transformation?

Generative AI and Scholarly Publishing

Tracy Bergstrom
Dylan Ruediger

October 30, 2024



ITHAKA S+R

Ithaka S+R provides research and strategic guidance to help the academic and cultural communities serve the public good and navigate economic, demographic, and technological change.

Ithaka S+R is part of ITHAKA, a nonprofit with a mission to improve access to knowledge and education for people around the world. We believe education is key to the wellbeing of individuals and society, and we work to make it more effective and affordable.

Copyright 2024 ITHAKA. This work is licensed under a Creative Commons Attribution Non-Commercial 4.0 International License. To view a copy of the license, please see <https://creativecommons.org/licenses/by-nc/4.0/>.



Introduction

In January 2024, Ithaka S+R published “The Second Digital Transformation of Scholarly Publishing” on the state of the scholarly publishing industry as it navigated “the second digital transformation” and assessed shared infrastructure needs in light of ongoing change in the structures, workflows, incentives, and outputs in the industry.¹ That report was based on interviews conducted in the first half of 2023 and included brief references to generative AI, which was just beginning to make its presence felt in academia and society. Since then, generative AI has become inescapable. As a tool that is capable of generating content, its implications for how scholarly research is conducted and for scholarly publishing and communication are potentially transformative.

Generative AI has already established a foothold in the industry: recent estimates suggest that perhaps one percent of the scholarly literature produced in 2023 shows signs of having been created in part with the assistance of a large language model (LLM).² Major publishers and content aggregators have rapidly developed and released AI-enhanced search and discovery tools and, less visibly, are experimenting with its potential use in back-end processes. The stage seems set for exponential growth in its use across the research and publication lifecycle.

What is not yet clear is how disruptive this growth will be. To this end, we interviewed 12 leaders in stakeholder communities ranging from large publishers and technology disruptors to academic librarians and scholars. The consensus among the individuals with whom we spoke is that generative AI will enable efficiency gains across the publication process. Writing, reviewing, editing, and discovery will all become easier and faster. Both scholarly publishing and scientific discovery in turn will likely accelerate as a result of AI-enhanced research methods. From that shared premise, two distinct categories of change emerged from our interviews. In the first and most commonly described future, the efficiency gains made publishing function better but did not fundamentally alter its dynamics or

¹ Tracy Bergstrom, Oya Y. Rieger, and Roger Schonfeld, “The Second Digital Transformation of Scholarly Publishing,” *Ithaka S+R*, January 29, 2024, <https://doi.org/10.18665/sr.320210>.

² Chris Stokel-Walker, “AI Chatbots Have Thoroughly Infiltrated Scientific Publishing,” *Scientific American*, May 1, 2024, <https://www.scientificamerican.com/article/chatbots-have-thoroughly-infiltrated-scientific-publishing/>.

purpose. In the second, much hazier scenario, generative AI created a transformative wave that could dwarf the impacts of either the first or second digital transformations.

These two scenarios are neither mutually exclusive, nor are they the only possible futures for generative AI.³ If generative AI proves to be genuinely transformative, it will presumably also create significant efficiency gains. As detailed below, certain aspects of scholarly publishing seem more likely than others to see significant change, leaving open the possibility that generative AI creates incremental change in some areas while disrupting others. Still, the scenarios are useful heuristic aids and were repeatedly used by our interviewees to frame their remarks and as indicators of their general disposition towards the technology itself. The strategic implications of generative AI for the publishing sector looks different from an incrementalist perspective than from a disruptive one, and even more complex if these different types of implications occur together.

Methodology

The observations in this report are based primarily on 12 interviews conducted between May and July 2024 (see Appendix A). The interviewees included representatives from publishers and closely aligned organizations, librarians, scholarly societies, and funders. We intentionally sought out individuals with broad, but deeply informed, expertise on generative AI and key trends in the business of scholarly publishing and communication. The semi-structured interview guide that was used to structure our conversations is included in this report (see Appendix B).

Because this report is designed as a companion or addendum to the larger “Second Digital Transformation of Scholarly Publishing” report, we decided to closely follow its internal structure. This decision facilitates reading the two reports together. Akin to the initial report, we frequently refer to “publishing organizations,” within this report as inclusive of publishers, repository services, and related providers that offer publishing

³ For additional scenario planning, see “ARL/CNI AI Scenarios: AI-Influenced Futures,” *Association of Research Libraries, Coalition for Networked Information, and Stratus Inc.*, June 2024, <https://doi.org/10.29242/report.aiscenarios2024>.

services to disseminate scholarly content to broad audiences.

We thank STM Solutions and six of its member organizations who supported this project. Our colleagues Claire Baytas, Oya Rieger, and Roger Schonfeld also provided valuable feedback at key points in the process. We're also grateful to Gary Price, who generously shares news about generative AI and higher education with us on a daily basis. The research and analysis presented here belong solely to the authors, and we accept all responsibility for its findings and conclusions.

Strategic Context

In this section, we provide an outline as to the primary strategic contexts that scholarly publishing currently faces, and how generative AI fits within these contexts. As in “The Second Digital Transformation of Scholarly Publishing,” we focus on the tremendous opportunity—and concomitant uncertainties, challenges, and diverging perspectives—that generative AI presents.

Transitioning towards Service Provision

As we noted in our earlier report, scholarly publishing as a whole is in the midst of a long-term shift away from a model centered on editorial work towards one based on services and platforms. We expect generative AI to accelerate this trend: publishing organizations are already engaged in strategic planning about how to map generative AI services to support the workflows of readers, authors, and editorial staff.

The boundaries between discovery, interpretation, and writing practices have already become increasingly interconnected. In the near future, we anticipate that discovery, interpretation, and writing services will become a fully integrated suite of tools aimed at keeping researchers engaged with a single platform across the research process. The rapid growth of tools like Digital Science's Dimensions AI Assistant and Clarivate's Web of Science Research Assistant are expanding the meaning of search and discovery through summarization or extraction features, and chatbot style interfaces

that allow users to query individual documents or whole corpora.⁴ We are also seeing experimentation with tools such as Paperpal, Writefull, or Curie that are designed specifically to improve the quality of academic writing.⁵

No one has yet combined discovery, interpretation and writing tools into a single product, but generative AI is creating new opportunities for publishing organizations to offer a wide range of author services on their platforms. Looking even further ahead, one can imagine hybrid discovery, summarization, and writing tools merging or becoming interoperable with enterprise research information management and analytics systems. In such a scenario, which would presumably require resources that only the largest actors in the industry can marshal, publishing organizations could see their core business transformed into being multipurpose providers of comprehensive research infrastructure.

Most of the individuals we interviewed believed that search and discovery will be heavily impacted by generative AI. Currently, researchers progress within search interfaces from discovery to understanding, but generative AI could fundamentally disrupt this linear progression with new incremental steps of generative AI-enabled synthesis. But many of the individuals we interviewed described the effects of generative AI on search and discovery more modestly, seeing it as changing the pace of scholarly communication and discovery rather than catalyzing substantive changes to the nature of research and knowledge creation.

The peer review process is a second area where generative AI is poised to make an impact. Peer review is a notoriously stressed component of the scholarly publication process, and most individuals we interviewed expressed some degree of optimism that generative AI will mitigate challenges associated with identifying potential reviewers and reduce the workload of reviewers and editors. “Publishers are working very hard to see if AI can do high quality peer review,” noted one interviewee, “they see this as their biggest bottleneck and hope they can find a way to solve the issue with Gen AI.”

⁴ See: Dimensions: A Digital AI Solution, <https://www.dimensions.ai/products/artificial-intelligence/>; and “Clarivate Launches Generative AI-Powered Web of Science Research Assistant,” *Clarivate*, September 4, 2024, <https://clarivate.com/news/clarivate-launches-generative-ai-powered-web-of-science-research-assistant/>.

⁵ See: Paperpal, <https://paperpal.com/homev2/>; Writefull, <https://www.writefull.com/>; and Curie, <https://www.aje.com/curie/>.

Specific ideas varied as to how generative AI may facilitate peer review in the near future.⁶ One person thought that it could help provide what was essentially “pre-review” feedback for human reviewers that would identify significant problems or help researchers understand if their work was a good fit for a specific journal. Others envisioned generative AI as a participant in peer review, with its contributions focused on issues such as making sure the abstract matched the text of the article, copy edits, and detection of research misconduct—including the undisclosed use of generative AI. At least one mentioned the possibility that generative AI could help publishing organizations more efficiently identify reviewers.⁷

Incorporating generative AI into peer review would presumably accelerate the flow of scholarly communication, while reducing friction points that burden authors, reviewers, and editorial staff alike. In the process, scholarly publishers would be better able to serve the scholarly community and advance scientific and creative discovery. There are also financial gains to be made by speeding up and streamlining the review process. Indeed, several of our interviewees believed that perhaps the most financially valuable application of generative AI for publishing organizations was its potential in this area. But the appeal of AI-enhanced peer review is driven by a mix of economic and mission motivations.

Even the most optimistic advocates for AI’s potential in peer review recognize the need for careful consideration of how to use generative AI to enhance rather than substitute for human engagement and knowledge throughout the review process.

However, the introduction of generative AI into the review process would come with real risk. Generative AI’s accuracy is often poor and at best too inconsistent to be trusted with even modest responsibilities for peer review. It can also cut against the core ethos of review by peers, a process

⁶ For recent discussion around this issue, see also Meadows et al., “Peer Review Week 2024: Ask the Chefs,” *The Scholarly Kitchen*, September 20, 2024, <https://scholarlykitchen.sspnet.org/2024/09/20/peer-review-week-2024-ask-the-chefs/>.

⁷ See, for instance, the functionality of Prophy: <https://www.prophy.ai/solutions/scientific-publishers/>.

that stresses the importance of expert, human, judgment by scholars with deep understandings of the context of a research output. Even the most optimistic advocates for AI's potential in peer review recognize the need for careful consideration of how to use generative AI to enhance rather than substitute for human engagement and knowledge throughout the review process.

Publishing organizations are cognizant of these risks. Concerns about confidentiality breaches associated with reviewers uploading manuscripts to LLMs have led some publishers to prohibit all use of generative AI.⁸ Others do allow limited use by reviewers to “polish, condense, or otherwise lightly edit” their reviewer reports, while retaining prohibitions against using it for other aspects of the review process.⁹ Funders are making similar decisions around the use of generative AI by grant reviewers.¹⁰

Confidentiality concerns could be mitigated by secure peer review environments where manuscripts and reviewer reports would not be used for model training. Microsoft's enterprise version of Co-pilot and enterprise versions of ChatGPT, both of which are being adopted by a growing number of colleges and universities, offer secure environments, as does

⁸ “The Use of LLMs or AI Tools in Peer Review,” Sage, <https://us.sagepub.com/en-us/nam/using-ai-in-peer-review-and-publishing#pt3>; Annette Flanagan, Jacob Kendall-Taylor, Kirsten Bibbins-Domingo, “Guidance for Authors, Peer Reviewers, and Editors on Use of AI, Language Models, and Chatbots,” *JAMA* 8 (2023): 702–703, <https://jamanetwork.com/journals/jama/fullarticle/2807956>; “Instructions for Peer Reviewers,” Cambridge University Press, <https://www.cambridge.org/core/journals/flow/information/peer-review-information/instructions-for-peer-reviewers>.

⁹ “Appropriate Use of AI-Based Writing Tools,” American Physical Society, 2024, <https://journals.aps.org/authors/ai-based-writing-tools>; “AI Policy,” Taylor & Francis, 2024, <https://taylorandfrancis.com/our-policies/ai-policy/#:~:text=Peer%20reviewers%20are%20chosen%20experts.the%20creation%20of%20their%20reviews>; “Wiley Peer Review Policy,” Wiley, 2024, <https://authorservices.wiley.com/Reviewers/journal-reviewers/tools-and-resources/review-confidentiality-policy.html>; “Generative AI Policies for Journals,” Elsevier, 2024, <https://www.elsevier.com/about/policies-and-standards/generative-ai-policies-for-journals>.

¹⁰ “Use of Generative Artificial Intelligence in Application Preparation and Assessment,” UK Research and Innovation, September 23, 2024, <https://www.ukri.org/publications/generative-artificial-intelligence-in-application-and-assessment-policy/use-of-generative-artificial-intelligence-in-application-preparation-and-assessment/#section-our-policy>; “Generative AI Tools Merit Caution,” National Institute of Allergy and Infectious Diseases, December 20, 2023, <https://www.niaid.nih.gov/grants-contracts/nih-case-study-copy-paste#:~:text=Generative%20AI%20Tools%20Merit%20Caution.the%20NIH%20Peer%20Review%20Process>; “Notice to Research Community: Use of Generative Artificial Intelligence Technology in the NSF Merit Review Process,” National Science Foundation, December 14, 2023, <https://new.nsf.gov/news/notice-to-the-research-community-on-ai>.

Amazon’s Bedrock.¹¹ Publishers and editorial systems providers will need to build this functionality into their tools and platforms. Overall, our interviewees consistently indicated that peer review would be among the aspects of scholarly publishing that are most likely to be impacted by generative AI in the short term. These initial impacts could yield significant efficiency gains that would create competitive advantages for companies that deploy them most effectively.

However, strategic planning should take more radical outcomes into account to address both efficiencies and innovation. Several interviewees expressed concern that publishing organizations were not thinking expansively enough about how to fully leverage generative AI. As one individual put it, publishing organizations should be thinking about generative AI as a technology that can and should allow us to do things that could not previously be done instead of thinking about it as a tool that lets us replicate what we already do more efficiently. The potential for platforms combining publication services with research information and analytic tools described above are one indication of those possibilities.

Consolidation and Competition

For several decades, consolidation has been one of the major trends in the scholarly publishing industry, with the five largest academic publishers now controlling over 60 percent of the market for journal articles.¹² Competition over the platform, analytics, and author services businesses in the sector has been equally fierce. Generative AI has implications for each of these business lines. As several of our interviewees noted, despite the widespread expectations that generative AI will create new revenue streams and affect business lines, how exactly it will do that is not yet clear.

¹¹ “Introducing ChatGPT Edu: An Affordable Offering for Universities to Responsibly Bring AI to Campus,” OpenAI, May 30, 2024, <https://openai.com/index/introducing-chatgpt-edu/>; Microsoft 365 Copilot, <https://www.microsoft.com/en-us/microsoft-365/copilot/enterprise>; Amazon Bedrock, <https://aws.amazon.com/bedrock/>.

¹² David Crotty, “Quantifying Consolidation in the Scholarly Journals Market,” *The Scholarly Kitchen*, October 30, 2023, <https://scholarlykitchen.sspnet.org/2023/10/30/quantifying-consolidation-in-the-scholarly-journals-market/>; Simon van Bellen, Juan Pablo Alperin, Vincent Larivière, “The Oligopoly of Academic Publishers Persists in Exclusive Database,” arXiv [preprint], June 25, 2024, <https://doi.org/10.48550/arXiv.2406.17893>.

Licensing content to large commercial LLMs, however, stands out as a significant exception to uncertainty about how to monetize generative AI. Several publishers, including Wiley and Taylor & Francis, have recently signed such agreements.¹³ These deals have attracted some controversy among scholars, but it seems reasonable to anticipate that other publishers will make similar licensing agreements in the near future. Ithaka S+R has launched a Scholarly Content LLM Licensing Tracker for those interested in more detail about these deals.¹⁴ Content licensing is a relatively indirect means of generating revenue from generative AI that does not require publishers to develop their own tools. It is also a use case that is minimally disruptive to an industry for which content licensing is a foundational part of the business model. It is difficult at present to imagine how smaller and midsize publishers will, if they are interested in doing so, license their content to LLMs, though there are a variety of intermediaries currently proposing to aggregate content in order to broker these deals. For larger publishers, decisions about whether to license content comes with strategic risk including whether the economic value of subscriptions will be undermined by making scholarly content accessible through an LLM, and whether the very idea of a version of record, a concept at the core of many publisher's business models, is diminished if their content is increasingly accessed indirectly through machine generated summaries.

Generative AI has also injected new competition in the search and discovery space. Tools designed to help scholars locate and understand relevant information, for example, Elsevier's Scopus AI and Consensus have proliferated over the last 12 months.¹⁵ Generative AI's summarizing and synthesizing capabilities were seen by most interviewees as one of its most far-reaching impacts. As one interviewee noted, the use of LLMs to make search more effective is not, in and of itself, a "huge game changer," but tools that can extend and elevate the research process beyond discovery could be in the near future. Many of the early AI

¹³ Christa Dutton, "Two Major Academic Publishers Signed Deals with AI Companies. Some Professors Are Outraged," *The Chronicle of Higher Education*, July 29, 2024, <https://www.chronicle.com/article/two-major-academic-publishers-signed-deals-with-ai-companies-some-professors-are-outraged>; Kathryn Palmer, "Taylor & Francis AI Deal Sets 'Worrying Precedent' for Academic Publishing," *Inside Higher Ed*, July 29, 2024, <https://www.insidehighered.com/news/faculty-issues/research/2024/07/29/taylor-francis-ai-deal-sets-worrying-precedent>.

¹⁴ "Generative AI Licensing Agreement Tracker," Ithaka S+R, <https://sr.ithaka.org/our-work/generative-ai-licensing-agreement-tracker/>.

¹⁵ See the description of tool features in Ithaka's S+R's Generative AI Product Tracker: <https://sr.ithaka.org/our-work/generative-ai-product-tracker/>.

enhanced search and discovery tools are based on abstracts and metadata rather than the full text of articles, but publishers and aggregators are moving towards integrating generative AI summarization tools into core content collections. At the moment, these tools are often marketed as premium features. However, we anticipate that some generative AI functionalities will soon become a standard component of the search and discovery capabilities of large scholarly collections such as ScienceDirect or JSTOR.¹⁶

The value of search and discovery tools is determined in large part by the content that those tools can access, and publishers have large and ever-growing collections of high-quality content behind paywalls that will be part of the value proposition of AI tools they create. However, the growth of open access publication has created multiple corpora of scholarly content that circulate freely on the internet. This corpora is available to independent developers to use and, presumably, has already been ingested by ChatGPT, Llama, Claude, and other foundational LLMs.

The legality of using scholarly and other forms of copyrighted content to train LLMs is contentious. *The New York Times* has filed one of several lawsuits alleging that using their content to train LLMs is a violation of copyrights.¹⁷ For its part, OpenAI argues that its use of “publicly available internet materials is fair use, as supported by long-standing and widely accepted precedents.”¹⁸ However this and similar cases are resolved, use of content published under some CC licenses would seem to on their face to prohibit or circumscribe their use by LLMs. CC-BY licenses, for example, require attribution as a condition of use, but the way generative AI uses sources to generate outputs often makes attribution nearly impossible. At least some publishers believe this means that training without permission on materials with these licenses is forbidden. Other CC license types prohibit commercial reuse entirely. Large tech firms may also be ingesting paywalled content. Sage has recently alleged that they have evidence that

¹⁶ Ithaka S+R and JSTOR are both services of the nonprofit ITHAKA.

¹⁷ Michael M. Grynbaum and Ryan Mac, “The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work,” *New York Times*, December 27, 2023, <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html>.

¹⁸ “OpenAI and Journalism,” OpenAI, January 8, 2024, <https://openai.com/index/openai-and-journalism/>.

tech companies have already and illegally “harvested much of our content to train large language models.”¹⁹

It will be years before the issue of when generative AI is transformative fair use, or a violation of copyrights or licenses, is established through litigation. In the meantime, publishers and aggregators will need to compete with nimble start ups and massive tech firms directly. As one individual noted, large publishers are feeling pressure from both sides and have responded by rushing products to market in an attempt to stake their claims.

For the most part, interviewees worried more about competition from big tech companies than from start-ups. They described several types of risk. One is what several interviewees described as the Google Scholar scenario, in which commercial LLMs could essentially outflank publishers, building out platforms and tools (or providing them through GPT store) that would become scholars’ default way of interacting with the scholarly record. The deals that many universities are signing with Microsoft or OpenAI to provide sandboxed research environments for scholars could, for example, acclimate researchers to using CoPilot or ChatGPT for research purposes and as an intermediary to accessing the scholarly record. This scenario would presumably not alter researchers’ desire to publish in scholarly journals, but it would complicate efforts to build or expand services across the research lifecycle. It might also, as one interviewee described it, encourage conservatism within the industry as publishers and aggregators seek to “create walled gardens” to buttress the exclusivity of their content, and miss opportunities to “recognize the fact that permissible use of content will add value.”

Another risk is that the use of generative AI to summarize or synthesize scholarly outputs leads fewer researchers to engage directly with articles, setting off a decline in readership, and a corresponding decline in clicks and other metrics used to measure the value of publisher and aggregator collections. As discussed in greater detail below, we heard quite a bit about the urgent need to build metrics that could account for the probability that readers would interact with scholarly publications primarily through the intermediary of a ChatBot.

¹⁹ Matilda Betterby, “Sage Confirms it is in Talks to License Content to AI Firms,” *The Bookseller*, September 19, 2024, <https://www.thebookseller.com/news/sage-confirms-it-is-in-talks-to-license-content-to-ai-firms>.

In our interviews, the “Google Scholar scenario” was evoked primarily to describe situations that would put scholarly publishing, or at least its current business models, at grave risk of losing touch with its audience. However, Google and Google Scholar also serve as an excellent example of the complexity, unevenness, and unpredictability of change. Many observers believed that Google would either kill off the academic library or that libraries would be completely immune from change. In fact, the development of Google’s search tools for the scholarly literature foreclosed a number of paths that libraries or publishers might otherwise have pursued but without killing either. Depending on how LLMs develop, and how publishers respond and lead, there may be a variety of different directions ahead.²⁰

A final risk identified by several interviewees was that the sector would underestimate how disruptive generative AI will be, and how rapidly those disruptions will unfold. As one person put it, the growth of ChatGPT is happening much more rapidly than adoption of earlier technological disruptors like Google. They emphasized that publishing organizations needed to be planning now for change so dramatic that even Google itself may no longer exist within three to five years. Publishing organizations would need to innovate to avoid becoming obsolete.

Humans and Machines

“The entire value chain of scholarly publishing is language based,” said one interviewee, “and thus generative AI can affect every aspect of publishing.” By creating new opportunities for human/machine interaction across that value chain, generative AI raises urgent questions about where, when, and how, human labor and knowledge add essential value to scholarly communication. These questions have complex ethical, practical, and legal components. Generative AI opens up a new phase of scholarship in which a human researcher may be the respondent, rather than the instigator, of new avenues of inquiry. Guardrails around usage, for which we become confident we share an understanding, are therefore imperative.

²⁰ Roger C. Schonfeld, “Tracking the Licensing of Scholarly Content to LLMs,” *The Scholarly Kitchen*, October 15, 2024, <https://scholarlykitchen.sspnet.org/2024/10/15/licensing-scholarly-content-llms/>.

“The entire value chain of scholarly publishing is language based...and thus generative AI can affect every aspect of publishing.”

Publishers have implemented initial policies around the use of generative AI by authors. These typically revolve around disclosure of use, a reminder that authorship cannot be ascribed to AI and may include language distinguishing use cases that are permitted (such as copy-editing) and those that are not. Taylor & Francis, for example, advises authors not to submit manuscripts where “generative AI tools have been used in ways that replace core researcher and author responsibilities,”²¹ while Elsevier prohibits use of generative AI to “replace key authoring tasks such as producing scientific, pedagogic, or medical insights, drawing scientific conclusions, or providing clinical recommendations.”²² These are valuable distinctions, but they are unlikely to hold for long, as generative AI will be used in ways that elide them.

The challenge here goes beyond disclosure: we will need a new conceptual understanding and vocabulary to capture the nuances of human/machine collaboration that generative AI is likely to enable. Several interviewees described this as an area where collective action by publishers would be useful as it would ensure a level of consistency across the industry. Others emphasized the limits of publishing organizations in this area, noting that it will ultimately be the scientific community that decides what uses of generative AI are acceptable. This will be challenging, as norms around the use of generative AI will take time to develop and meanwhile early adopters and power users will continue to find new uses for it.

One additional challenge that will require coordinated and careful consideration is how increased use of LLMs by researchers can be

²¹ “AI Policy,” *Taylor & Francis*, 2024, <https://taylorandfrancis.com/our-policies/ai-policy/#:~:text=Authors%20should%20not%20submit%20manuscripts.code%20generation%20without%20rigorous%20revision>.

²² “The Use of Generative AI and AI-Assisted Technologies in Writing for Elsevier: Policy for Book and Commissioned Content Authors,” *Elsevier*, 2024, <https://www.elsevier.com/about/policies-and-standards/the-use-of-generative-ai-and-ai-assisted-technologies-in-writing-for-elsevier>. For an overview, see “Generative AI in Scholarly Communications: Ethical and Practical Guidelines for the Use of Generative AI in the Publication Process,” a white paper published in December 2023 by STM Solutions.

balanced with the imperative to uphold the fundamental tenets of scholarship, including respect for provenance, attribution, reproducibility, and transparency. LLMs can already create, or assist in the creation, of new knowledge, but they lack effective safeguards to ensure researchers are understanding the source material or provenance of this content. Additionally, new scholarship created with LLMs must be transparent and reproducible in the ways in which we would expect of traditional forms of scholarship. LLMs do not currently facilitate this interplay that is so critical to the proliferation of creating new knowledge.

Adapting to generative AI also will require decisions about where human judgment is necessary in conducting research and writing articles, in peer review, and throughout the editorial process. As one interviewee put it, the industry and its stakeholder communities are ultimately going to invert the questions of what role machines can play in favor of asking what role humans need to play. There seems to be some consensus that humans will need, at least, to validate and review AI generated content, but specifics about how to preserve space for human cognition and understanding often consisted of abstract principles about keeping “humans in the loop.”

Publishing organizations recognize that they will need to err on the side of preserving roles for humans and that they will meet resistance if they move too quickly. One of the paradoxes of the moment is that the caution with which publishing organizations are likely to proceed cuts directly against the need to innovate rapidly. However, keeping human understanding at the center of the scientific record may face challenges from several directions.

Commercial forces are one such challenge, and several interviewees expressed fears that market pressures would ultimately lead publishing organizations to a slippery slope of ever more aggressive automation. The temptation to increase the speed of the publication process and the pull of revenue and/or efficiency gains, one said, is going to be very hard to resist. Academic reward structures are another source of pressure, as both authors and peer reviewers will also be incentivized to cut corners, creating the potential for a downward spiral that could cause lasting damage to the trustworthiness of the scholarly record.

Finally, the dynamics of machine learning itself could diminish the fundamental purpose of scholarly communication by marginalizing human readers. Generative AI will most clearly affect what might be considered

the “front” end of scholarly publishing (discovering and reading sources) and the “back” end of writing, editing, and reviewing articles documenting the experimental process. The labor that is expended on these tasks is at the core of scholarly publishing, but its value is predicated on the need for human readership and interpretation of the scientific record to advance discovery. As we have described above, the synthesizing, extracting, and summarizing capabilities of generative AI create a new type of intermediary between readers and the scholarly literature. If AI becomes the primary reader of this content, it raises the question of whether the work of creating a human readable scientific record will come to be seen as an unnecessary expense—or as a barrier to improved machine readability.

Looking even further ahead, frameworks for automation will span across the entire research process from hypothesis generation to experimentation and creation of journal articles.²³ Whether these frameworks are viable is unclear, but they illustrate just how marginal the humans in the loop could potentially become and how fundamentally the communicative purpose at the heart of scholarly publishing could change.

Establishing a Trusted Global Public Good

While scholarship should be considered a trusted global public good, fraud, misconduct, and other malevolent activities of the past decade have challenged this idea. Generative AI, and its potential for producing falsified content across media, has pushed the temperature of longstanding conversations about maintaining trust and trustworthiness in scholarly publishing to a new level. While the risks posed by generative AI to trust are evident, many of the people we spoke with were optimistic about its potential to make scholarly publishing more accessible to scholars and readers around the globe, and more useful as a public good.

Generative AI was widely viewed as leveling the playing field for authors and readers. There is a perception that researchers for whom English is not a first language are already making use of generative AI to improve the quality of their academic writing. Indeed, we heard reports that vendors

²³ Eliza Strickland, “Will the ‘AI Scientist’ Bring Anything to Science?: A Tool to Take Over the Scientific Process Continues a Controversial Trend,” *IEEE Spectrum*, September 9, 2024, <https://spectrum.ieee.org/ai-for-science-2>.

who supply copyediting services targeted at non-native speakers are seeing steep revenue declines and could become casualties of generative AI. For the most part, though, our interviewees were excited by this development, which was seen as advancing equity and access to journals published in English. As one person noted, “English is the mode for communicating science for many people right now,” and generative AI can “level the playing field for all researchers to communicate in this space.” Beyond its value as a tool for equity, the person said, it also contributes to “a world where the best science wins out.”

Interest was similarly widespread in the possibility that translation could be largely automated and thus scaled. In theory, generative AI could jump start a future where the entire scholarly record is accessible to speakers of major languages. In the process, the consumer base for the content produced by scholarly publishing organizations would grow dramatically, opening the door to greater globalization of the market.

Calculating Impact

The second digital transformation created new standards such as citation indices for assessing the impact of individual researchers and the value of publisher’s offerings in the form of COUNTER metrics. Both metrics may be profoundly disrupted by generative AI if scholars come to rely on using generative AI as an intermediary method for using the scholarly record.

Our interviews surfaced a particularly acute need for developing metrics to enhance COUNTER metrics, which provide critical common ground for libraries and publishing organizations regarding the value of their collections.²⁴ We encountered little in the way of concrete suggestions about how to measure engagement. Two primary challenges emerged.

The first is that traditional COUNTER metrics provide counts of engagement with items that include a Unique Resource Identifier (URI). While this may include videos, metadata, articles, book chapters, or other content, it is imperative that the resource have a distinct identifier and be discoverable by multiple researchers to be counted. Generative AI,

²⁴ See a review of related issues in Tim Lloyd, “Guest Post: Time to Rethink Usage Analysis,” *Scholarly Kitchen*, October 2, 2024, <https://scholarlykitchen.sspnet.org/2024/10/02/guest-post-time-to-rethink-usage-analytics/>.

however, encourages the creation and consumption of bespoke content as generated by the individual researcher, for example a literature review generated in response to a specific topic. This resulting content is ephemeral in that it is generated for the individual and no formal record remains once the researcher's inquiry is concluded. With these parameters, this use case at present would not be counted as an investigation or readership of the included content.

The second issue is that traditional metrics do not measure the level of engagement with a particular resource. COUNTER metrics' traditional division of investigations versus requests, for example, assess a researcher's engagement with supplementary versus full-text content. While these counts are meaningful in evaluating what content has been accessed and by what means, there is no measurement of prolonged engagement with a resource. As different Generative AI tools make more granular encounters with content possible, including that researchers will be able to make multiple, adaptive inquiries to the same article or resource, future metrics will ideally allow for additional specificity in understanding the depth at which researchers are responding to and engaging with individual resources.

New Opportunities for the Shared Infrastructure

In this section, we build upon the strategic context examined above to discuss opportunities to create new categories of shared infrastructure. Because the future development of generative AI is highly speculative, we have focused on areas of opportunity for development or strengthening of the generative AI landscape as it relates to scholarly publishing and its usage within the generation or consumption of the scholarly record.

Spine of the Scholarly Record

The “Second Digital Transformation of Scholarly Publishing” report contained a section devoted to the “spine of the scholarly record,” which we defined as a new kind of standard and the associated infrastructure needed to ensure that the scholarly record can be effectively organized and maintained. In our contemporary research environment, in which an article may be atomized into component parts including a preprint, an associated dataset, related code, or other elements, it is unclear how these components may be persistently linked and preserved to ensure future usability. New infrastructure is needed to extend the notion of the version of record and to provide a spine around which all of these associated research outputs are organized.

The advent of LLMs, which generate output and summaries in ways that can make their source material difficult to cite or understand, will make it more difficult for researchers to understand the provenance and context of the information they are receiving. This remains true despite upgrades that now allow LLMs to include citations in its responses (though not always citations to actual sources). Emerging research suggests that LLMs, even when supplemented by Retrieval Augmented Generation (RAG), often provide unsupported citations and use them in ways that strip them of critical context.²⁵ A related challenge is that LLMs, by design, draw on and transform information in ways that do not lend themselves to clear citation and make the provenance of responses very difficult to trace back to its initial context.

Uncovering a dataset is useful, for example, but potentially problematic if it has lost the context of its broader study, information about its source of funding, and outcomes. In this sense, however, the advent of big data and automated data aggregation helps us collectively to understand what data is and its inherent limitations. Science and data are not equivalent; as one interviewee stated, “science is the narrative through which the data is explained.” This narrative is structured differently depending on the field. One area in which generative AI could be advantageous in the near term is

²⁵ See especially Kevin Wu et al., “How Well Do LLMs Cite Relevant Medical References? An Evaluation Framework and Analysis,” arXiv [preprint], February 3, 2024, <https://arxiv.org/abs/2402.02008>; Courtni Byun, Piper Vasicek, and Kevin Seppi, “This Reference Does Not Exist: An Exploration of LLM Citation Accuracy and Relevance,” *Proceedings of the Third Workshop on Bridging Human-Computer Interaction and Natural Language Processing*, June 21, 2024, <https://aclanthology.org/2024.hcinlp-1.3.pdf>.

to bridge gaps between disciplines through automated and stronger documentation of data that could allow data to be more widely shared and utilized.

There is underlying risk here to publishing organizations, however, if publishing becomes more data-centered—as the data becomes more central with the article potentially no longer in the center of the network of atomized components—publishers’ roles may lose critical value as the narrative as expressed in an article loses centrality.



Recommendations

- Publishing organizations and content aggregators should prepare for scenarios in which users’ primary interactions with scholarly outputs are mediated by generative AI. We recommend that publishing organizations and stakeholders collaborate to create standardized metadata or other provenance markers including interconnected persistent IDs that can travel with snippets or synopses of scholarly outputs to facilitate citability, transparency, impact metrics, and improve the quality of generative AI outputs.
- Research communities need to build consensus about when prompts, responses to prompts, and other forms of generative AI generated content or queries should be recognized as part of scholarly outputs and, if so, how they should be cited. We recommend that cultural heritage organizations, libraries, and content aggregators work to build consensus about the historical value of these new content forms and develop frameworks for prioritizing preservation.

Research Integrity

One overriding theme we heard across interviews was that the quality of content that currently underlies LLMs is not reliable enough to ensure the integrity of the scholarly record. This issue of the lack of trustworthiness of content manifests in several different ways. At the time of writing this report, several large publishers had recently disclosed licensing

agreements for their journal and/ or book data to be ingested into generative AI tools. Even with access to high quality data, LLMs make frequent errors of fact and interpretation. Without transparency about a model's underlying training data and clear understanding of how its responses are generated, researchers cannot be confident in the scholarly accuracy of their results.

As time passes, the advent of Small Language Models (SLMs) that rely on smaller but more focused training data, for instance with topical focus on specific academic discipline, may fulfill this need. Vendors of generative AI products for research often make use of Retrieval Augmented Generation (RAG), which grounds generative AI outputs in a knowledge base outside of the training data, to improve the accuracy and relevance of outputs. For example, both Elsevier's Scopus AI and Clarivate's Academic AI Platform point to RAG architecture as a means to assure their users of the quality and reliability of their products for research purposes.²⁶ As RAG technology continues to be improved, this will ideally ensure greater scholarly accuracy in generative AI-assisted research.²⁷

To ensure research integrity in conjunction with increasing usage of generative AI tools, interviewees frequently cited the need for new standards to ensure consistency and transparency. Building tools that allow for the understanding of the provenance of included information will be critical, as a researcher must be able to verify and vet the source of information on which new scholarship is based. We also frequently heard that publishers must articulate with more granularity how it is acceptable to utilize generative AI tools, and what types of use cases are reliable or not, so that authors understand areas in which they need to be transparent about use.

²⁶ See "Scopus AI," Elsevier, <https://www.elsevier.com/products/scopus/scopus-ai#1-the-scopus-ai-difference>; and Guy Ben-Porat, "Introducing the Clarivate Academic AI Platform," *Clarivate Blog*, May 21, 2024, <https://clarivate.com/blog/introducing-the-clarivate-academic-ai-platform/>.

²⁷ On the remaining room for improvement in RAG, see Kyle Wiggers, "Why RAG Won't Solve Generative AI's Hallucination Problem," *TechCrunch*, May 4, 2024, <https://techcrunch.com/2024/05/04/why-rag-wont-solve-generative-ais-hallucination-problem/?guccounter=1>; Yizheng Huang and Jimmy Huang, "A Survey on Retrieval-Augmented Text Generation for Large Language Models," arXiv [preprint], August 23, 2024, <https://doi.org/10.48550/arXiv.2404.10981>, specifically the section on "Challenges and Future Directions" on page 25; Shailja Gupta, Rajesh Ranjan, and Surya Narayan Singh, "A Comprehensive Survey of Retrieval-Augmented Generation (RAG): Evolution, Current Landscape and Future Directions," arXiv [preprint], October 3, 2024, <https://doi.org/10.48550/arXiv.2410.12837>.

Research libraries are well positioned to engage in work around ensuring the transparency and verifiability of the products and provenance of scholarly communication. Especially as libraries grapple with new expectations of licensing resources that both promote usage of new tools while maintaining the trustworthiness of the scholarly record, librarians have much to offer publishing organizations in terms of observing emerging methods of conducting research with generative AI tools. Several interviewees cited concern that publishing organizations do not have the capacity to keep up with new tools as utilized by authors. Increased efforts by librarians in this area may bridge this gap and also advance conversations about how licensing must evolve to allow for developing research that utilizes generative AI.

Academic fraud is a serious problem, and generative AI technology enables fraud to be conducted faster and in more disquieting ways. In conducting research for “The Second Digital Transformation of Scholarly Publishing,” a number of interviewees shared with us concerns about a substantial growth in fraudulent content submissions generated with AI techniques, for which they felt unprepared due to the volume and complexity of detection. One particular challenge is that generative AI can create fraudulent content across media, including text, images, video, and data, necessitating a constant evolution of new detection techniques.

Much of the current discussion on research integrity is framed around protecting the scholarly record from falsified content and dishonorable actors, but while these occurrences are highly problematic, they are also relatively rare. As several interviewees suggested, there is also value in using generative AI tools to formulate stronger metadata within datasets for reuse, or to make manuscripts stronger through features such as language revision, literature review assistance, publication readiness checks, and the like. In this sense, there are opportunities to think proactively about how generative AI could make positive contributions to research integrity.

Finally, the issue of ensuring research integrity provokes the larger speculative question of how we define trust, or trustworthiness, in a world in which machine-to-machine communication increasingly generates more of the scholarly record. As one interviewee described, generative AI “exacerbates and speeds up” already existing issues surrounding the erosion of trust in research and the research enterprise. Another interviewee noted that organizations will have to be more transparent in

providing trust indicators to the public of the integrity of their material, which internally will require tightening up processes. In sum, ensuring trustworthiness is not an issue with a singular answer, but one that will require diligence and effort as generative AI tools continue to evolve.



Recommendations

- Publishers will need to strengthen their role in advocating for and publishing high quality content as issues around research integrity become even more complex within the generative AI landscape. Curation/provenance is going to become even more important to scholarly publishers as the overall quality of the web as a source of information declines. We recommend that publishing organizations work together with technology providers to establish trust markers that identify peer reviewed content as visible within LLMs.
- Responsibility for ensuring the trustworthiness of scholarly outputs cannot be left primarily to publishers. We recommend frank discussion between actors from the entire research lifecycle, including funders, university research officers, scholars, scholarly societies, and publishing organizations to create more equitable, comprehensive, and effective ways to ensure the integrity of scientific research and better distribute responsibilities for identifying and correcting serious misconduct. We recommend that a broad range of organizations and individuals participate in these discussions to ensure transparency.

Making Meaning

The scholarly record enables a variety of communities to make meaning and generate new scholarship from preceding research. What constitutes the scholarly and cultural record, however, and its boundaries, become an increasingly important question with the advent of LLMs. How will LLMs absorbing major publishers' monographic and journal corpora—but not necessarily online exhibits or other gray literature—affect the future outputs of research and fundamentally change scholarship? As several

interviewees indicated, a great amount of information is available that is not in the scholarly publishing corpus proper, that we need to be mindful about potentially losing cultural heritage and memory if meaning is increasingly derived from LLM content.

On the contrary, new meaning from the scholarly record can be derived in ways not possible before. Generative AI tools are highly flexible and intuitive, greatly lowering the barrier of expertise for usage. Whereas before training as a data scientist was necessary to query a corpus of information in similar ways, generative AI interfaces have democratized and simplified this process. Generative AI tools can also be trained with non-linguistic data, such as scientific images, so there is immense opportunity to work with content from across different modalities.²⁸

The underlying question of where the human will be necessary within this work, however, is critical to the transformational possibility of the tools. Generative AI may allow us to make a great deal more meaning out of scholarship than previously possible if human scientists no longer have the monopoly on initiating the narrative. If generative AI tools are soon creating a first draft of a scientific narrative, at that point the human researcher is the reviewer. We are also beginning to see radical frameworks for completely automated science which leverage generative AI to generate research ideas, write necessary code, conduct experiments, and create visualizations and written outputs to communicate findings emerge.²⁹ Such changes would mark a transition from AI-enhanced research to machine led research, setting the stage for truly transformative changes in how research is conducted.

Generative AI tools may also be transformative in how we think about reading, or the consumption of engaging with the scholarly record. Assisting researchers in drafting literature reviews by streamlining the processes of identifying and evaluating relevant literature is already a prominent and appealing use case for existing generative AI research

²⁸ Examples of multimodal models include LLaVA (<https://www.microsoft.com/en-us/research/project/llava-large-language-and-vision-assistant/>); or, for the medical domain specifically, the proof of concept for a multimodal version of Med-PaLM. Tao Tu et al., "Towards Generalist Biomedical AI," arXiv [preprint], July 26, 2023, <https://doi.org/10.48550/arXiv.2307.14334>.

²⁹ Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha, "The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery," arXiv [preprint], August 15, 2024, <https://doi.org/10.48550/arXiv.2408.06292>.

products.³⁰ It is assumed that such products will lead to less reading and evaluation of articles by researchers themselves. If generative AI tools are eventually the ones composing a substantial, bespoke literature review and forming conclusions about areas of opportunity for research within a topic, this creates a scenario in which humans are supervising the parameters of the review and reading the output, but not reading the individual papers or components. Interviewees expressed mixed opinions whether or not this generative AI-enhanced way of consuming scholarly content represents a useful tool for researchers or diminishes the critical thinking necessary to truly engage in scholarship.

Key to each of these issues is that, at present, generative AI-enabled research results are not good enough in that they are frequently incorrect or incomplete, and cannot handle the nuances, complexities, and details of the problems researchers are trying to solve. Even if hallucination rates drop significantly, persistent concerns about attribution, provenance, and transparency of content need clarification. High quality, open access content with licensing guardrails including those that require attribution may drive future development in this area that raises the quality of content.



Recommendations

- Different publishers have different expectations about generative work, which reduces researchers' ability to take their work to a different journal or venue. We recommend the development of common vocabulary and shared expectations around usage to promote researcher understanding and compliance.
- The evolution of impact metrics will be critical to our collective understanding of how generative AI is being used by researchers, and in our understanding of the evolution of academic readership. We recommend the funding of a study to investigate how COUNTER metrics should evolve to meet these needs.

³⁰ See, for example, the University of Iowa's list of recommended tools for AI-assisted literature reviews: "[AI-Assisted Literature Reviews](https://teach.its.uiowa.edu/news/2024/03/ai-assisted-literature-reviews)," <https://teach.its.uiowa.edu/news/2024/03/ai-assisted-literature-reviews>, which includes a few of the many tools aimed at this use case.

Supporting New Business Models

Innovation in the generative AI space has been rapid over the past 24 months, and the pace of iteration has outstripped publishing organizations' abilities to adapt underlying business models. Individuals we spoke to across the sector acknowledged that organizations need to invest more time in understanding generative AI technology more deeply because it has the potential to upend a number of underlying systems in the future. At the time of writing, we did not learn of concrete examples of instances in which publishing organizations were utilizing generative AI tools in ways that substantially increased revenue in regard to their backend processes. While licensing of content to LLMs seems poised to become a new source of revenue, generative AI tools are not yet embedded within the daily work of publishing organizations and therefore are not yielding major efficiencies. In particular, interviewees were sensitive that careful consideration needs to accompany any uses of generative AI tools that may replace human employees.

In parallel to how open access has disrupted the fundamental business model of publishing organizations over the past two decades, generative AI has the same potential to be transformative for those with the resources to best adapt and harness its potential. One challenge is that smaller publishing organizations may not have the human or technological capital to keep pace with their larger counterparts, potentially resulting in further consolidation within the industry. Large commercial publishing organizations that hold significant content within their purview will be able to create new functionality and attract users, simultaneously making it more difficult to operate independently within this space.

Some discrete services within the broader publishing landscape have already seen significant disruption, including copyediting and translational services. In these instances, providers are investing heavily in generative AI tools in an attempt to shore up their value proposition. There is direct economic impact to this sector of the industry already in decline of revenue.



Recommendations

- We heard from multiple perspectives that, due to the enormous proliferation of generative AI tools over the past 18 months, authors and publishing organizations do not have a common understanding of its opportunities. We recommend that stakeholders work together to build this shared understanding of the potential value and harms of generative AI to scholarly communication.
- Open on-demand translation services embedded within generative AI tools have the potential to be transformative in enabling globally connected research. We recommend that the quality of forthcoming services be carefully vetted for usage within academic contexts. On-demand translation services also raise questions about current discounting models for non-Anglophone countries that need careful consideration to ensure incentives to global participation remain strong.

Conclusion

In the “The Second Digital Transformation of Scholarly Publishing,” we noted that few individual infrastructure providers have achieved all of the key characteristics that stakeholders value: one that provides financial returns necessary to sustain the enterprise and support innovation, is sufficiently agile to navigate change, and would meet standards of trustworthiness required to support scholarship. Generative AI injects an unpredictable new dynamic into the ecosystem, but it is unlikely to change the importance of these basic characteristics. Our interviews showed relatively high levels of confidence that generative AI will expedite innovation and provide tools to make the publishing environment more agile and flexible. However, we also surfaced the interviewees’ deep concern that generative AI would undermine the trustworthiness of the scholarly record, and we heard repeated questions about exactly how and when generative AI tools will provide positive revenue streams.

We also encountered considerable anxiety about the capital requirements of innovation in the space, which will likely disadvantage small publishers relative to their larger peers, and force even the largest publishers to rely on big tech to fuel innovation. In our earlier report, we wrote about the importance of certain types of infrastructure in helping to ensure the competitiveness of the publishing marketplace. This imperative is only more clear when we consider the impacts of generative AI. As we have written in a number of sections of this issue brief, the investment in integrating these technologies or licensing content to their providers is not trivial and benefits tremendously from scale. The question is not just whether a few large publishing organizations will be able to survive, and only further differentiate themselves, through creative investments and innovation, but also whether the shared infrastructures on which smaller publishers rely absolutely will be able to innovate in order that these smaller publishers can keep up. It is absolutely essential that smaller publishers also find their voices, and find ways to make investments, in order to remain in the game.

As indicated within the introduction, we nonetheless heard a great deal of optimism as to how generative AI can advance the ideals of scholarly communication. Most interviewees at this time see a future in which generative AI becomes embedded in processes and workflows across the publication lifecycle, making positive contributions to existing goals, processes, and infrastructures for scholarly publishing. However, the path to this future depends on stakeholders taking action *now* to ensure that definitions around what constitutes ethical usage of generative AI is understood consistently across authors, readers, rights holders, publishers, and aggregators. Strategic optimism about generative AI is predicated on finding ways to ensure trust and trustworthiness in scholarship.

Trust is not the only risk posed by generative AI, though it is perhaps the most significant and certainly the most easy to understand. There are harder to define futures in which generative AI sets off a series of transformations that fundamentally restructure the industry's players and even its basic purpose. Many of these pivot on whether generative AI distances scholars from the outputs they create, review, and consume, and whether the scholarly record becomes essentially a type of background material or training data that scholars interact with primarily via AI-generated summaries, snippets, or synopses. There has already been speculation of a future in which scholarly communication occurs

largely between machines: generative AI could make this future more likely. Such a scenario could easily sideline publishers, their platforms, and their content. This would require significant reimagining of the value of publishing organizations to scholarly communication and recognition, perhaps after the fact, of the importance of human governance and control over AI.

Appendix A: List of Interviewees

- Oren Beit-Arie, Senior Vice President, Strategy & Innovation, Clarivate
- Lorcan Dempsey, Professor of Practice and Distinguished Practitioner in Residence, University of Washington, Information School
- Judson Dunham, Senior Director of Product Management, Elsevier
- Dave Flanagan, Senior Director, Generative AI Product Strategy, Wiley
- Darla P. Henderson, Director, FASEB Open Science and Research Integrity and Director, Publications, Federation of American Societies for Experimental Biology
- Kate Hertweck, Program Manager, Open Science, Chan Zuckerberg Initiative
- Cynthia Hudson Vitale, Associate Dean of Technology Strategy and Digital Services, Johns Hopkins University, Sheridan Libraries and University Museums
- Stuart Leitch, Chief Technology Officer, Silverchair
- Ian Mulvany, Chief Technology Officer, BMJ Group
- Avi Staiman, Founder and Chief Executive Officer, Academic Language Experts
- Thomas Suetterlin, Vice President AI, Springer Nature
- Jennifer Wright, Head, Research Integrity and Publication Ethics, Cambridge University Press

Appendix B: Interview Questions

1. How would you describe your level of familiarity and expertise with AI in general and with generative AI tools specifically?
2. Where are you already seeing the effects of generative AI in the scholarly publishing industry? (examples could include *internal processes* such as identifying peer reviewers, *reader facing innovations* such as enhanced search and discovery, or *author focused innovations* such as writing/revision tools, or *negative consequences* such as increased plagiarism, falsified data, or the spread of misinformation).
3. To the best of your knowledge, are your organization and/or others in the sector developing new applications for generative AI that you anticipate will become prominent or integral to processes over the next 12-24 months?
4. How “ready” is the existing shared technical and/or social infrastructure for scholarly publication/communication to support the widespread adoption of generative AI tools?
 - Which components of this infrastructure do you think will be most transformed by generative AI?
5. What do you see as the primary ethical, competitive, or market challenges posed by generative AI? What kinds of tools or infrastructure will we need to ensure ethical adoption?
 - What systems and structures will be necessary to balance the needs of authors, readers, rights holders, publishers, and aggregators?
6. What does the most optimistic part of yourself think are the greatest opportunities generative AI presents to better support the goals of scholarly publishers? What does your most pessimistic self see as the greatest risks? On balance, are you more optimistic or pessimistic?

7. What parts/processes of the publishing sector do you think will be changed the most by generative AI? Where do you anticipate the least amount of change?
8. Thinking broadly about the communities and stakeholders involved in scholarly communication, where do you see the most urgent need for coordinated action or consensus building?
9. Are there any issues that we haven't talked about already that you think need to be raised?