

June 12, 2023

Travis Hall
National Telecommunications and Information Administration
U.S. Department of Commerce
1401 Constitution Avenue NW, Room 4725
Washington, DC 20230

Re: STM Response to National Telecommunications and Information Administration (NTIA) hereby requests comments on Artificial Intelligence (“AI”) system accountability measures and policies (Docket No. 230407-0093)

Dear Dr. Hall:

[STM](#), the International Association of Scientific, Technical and Medical Publishers welcomes the opportunity to comment on the “important federal objective” of promoting “trustworthy artificial intelligence” as part of US national security, innovation, and commerce consistent with US values. Our publisher members are increasingly participating in AI activities from at least three perspectives, including: as key providers of high-quality content, context and data; as users of AI in both internal workflows and services for authors, editors, and reviewers and as users of AI in external tools and services; and as key stewards of trust and integrity in the scholarly record and in public discourse. It is in these contexts that we offer the following comments.

Trust is at the center of what STM and its members do, as our tagline “advancing trusted research” attests. STM’s members are at the forefront of digital innovation, providing stored and organized information, tagging and enriching content, and creating ontologies. As such, they play an indispensable role in making available high-quality, trusted information and knowledge – traditionally for human consumption, but increasingly also as input for AI models and applications.

Publishers are key providers of information and data on which AI is run. Relevant, high-quality input and training data for AI developers and systems form one of the key ingredients for high-quality, trustworthy, and ethical outputs. Providing this corpus of information in requested digital formats is a core expertise of publishers. By validating, normalizing, tagging and enriching content, delivering material in robust, interoperable, and globally consistent formats, and creating domain-specific ontologies, publishers ensure that information is a trustworthy high-quality input source with tremendous potential for use by AI systems across a broad range of applications. Regardless of the purpose of an AI tool, the accuracy of the scientific record maintained by science and academic publishers helps to ensure that machine learning has both depth and accuracy.

In addition, STM and its members are actively considering how AI can be developed and utilized in an ethical, accountable, and trustworthy manner. Publishers are driving innovation in this space as well as considering applications of AI to the research enterprise. For example, ‘smart science’ applications of AI could go beyond testing hypotheses against vast amounts of data to also creating new ones, developing new theories, exploring new connections.

Whilst embracing AI technologies to enhance operational efficiency and drive new discoveries, publishers remain steadfast in upholding the principle of rigorous human oversight, which is crucial in ensuring the highest standards of quality and integrity and safeguarding against misinformation.

Publishers continue to prioritize the role and expertise of human reviewers and ultimately take editorial and legal responsibility for the content they publish.

Being aware of the need to consider issues related to trustworthy and responsible AI, STM published a White Paper in May 2021¹ outlining best practice principles for an ethical, trustworthy, and human-centric AI that discusses both the potential of AI tools and systems and also its risks. One of the sections focuses on transparency and accountability, and we refer federal stakeholders to this White Paper for further information and context on our comments. We also refer you to an STM submission to the US Patent and Trademark Office in 2020² on the need for intellectual property protection for AI innovation, which is relevant to accountability and ensuring the quality of inputs, AI systems, and outputs.

Fundamentally, there can be no transparency or accountability without understanding the quality of inputs to AI tools and systems. To ensure an understanding of the quality of inputs, as well as enable the continued availability and effectiveness of these potential inputs, STM makes several high-level recommendations:

- First, a clear distinction must be made between raw data on the one hand, and structured data and copyrighted works on the other hand. For the latter, licensing, rather than reliance on fair use or other exceptions and limitations, should in most instances be the method of choice for enabling access and use as training data for AI. Respect for copyright will foster an accountable and trustworthy AI environment.
- Second, there should be clear provenance for inputs to promote the integrity of AI tools and the reliability of their outputs, which is a key added benefit of licensing. Visibility into the data sources will also enable users and auditors to better tackle bias, to ascertain that a model was trained on data collected with the consent of those involved, and to ensure legal and regulatory compliance. In the long term, this will lead to higher quality and more trustworthy AI systems, especially by encouraging the use of the highest quality final version of any articles, that is, the “Version of Record.” Use of the trusted and maintained Version of Record will support accountability, trust, and accuracy and reduce the potential for erroneous output that could undermine science, public health, and understanding.
- Finally, when applying AI in a context of scholarly communications, a record of inputs and outputs to the AI system should be maintained to ensure that the AI system and its outputs can be placed into a chain of evidence and results can be more easily reproduced, including references to scholarly works that have been used.

Accountability in research and science is a responsibility shared by all key stakeholders, including researchers, funders, policy makers, and publishers. We recommend that the federal government work with publishers and other stakeholders and rely on existing standards and practices to promote accountability as referenced in the background to this RFC. STM is committed to working with NTIA and other stakeholders to support the development of responsible AI tools and systems and to establish new accountability standards and adapt other standards where necessary.

¹ STM, “AI Ethics in Scholarly Communication: STM Best Practice Principles for Ethical, Trustworthy, and Human-centric AI,” April 2021. https://www.stm-assoc.org/2021_04_29_STM_AI_White_Paper_April2021.pdf

² “STM SUBMISSION TO UNITED STATES PATENT AND TRADE MARK OFFICE, DEPARTMENT OF COMMERCE REQUEST FOR COMMENTS, 84 FR 58141, pp 58141-58142: INTELLECTUAL PROPERTY (IP) PROTECTION FOR ARTIFICIAL INTELLIGENCE (AI) INNOVATION,” January 10, 2019. <https://www.uspto.gov/sites/default/files/documents/International%20Association%20of%20Scientific%20and%20Medical%20Publishers%20RFC-84-FR-58141.pdf>

STM also supports the views expressed by others in the scholarly publishing ecosystem as expressed in their replies to this RFC, particularly those of the Association of American Publishers (AAP) and the Copyright Clearance Center (CCC). In addition, we are providing comments on the following specific questions.

1a. What is the purpose of AI accountability mechanisms? What kinds of topics should AI accountability mechanisms cover?

AI accountability mechanisms should both enable end users to evaluate the validity and trustworthiness of AI tools, as well as ensure that AI tools are compliant with legal standards, regulations, and the claims of the AI tool and developer. A key topic therefore should be the clarity and transparency about data and works used by and in the development of the AI tool or system, as well as assurance of appropriate rights and licenses to use the inputs and outputs of the AI.

For the scholarly and research ecosystem, accountability is critically important to enable outputs to be relied upon, especially when scholarly literature is used by AI tools or systems to generate any form of insight, recommendation, or knowledge (e.g., summary, layman version, answer to a specific question, etc.) based on the literature. The scholarly discourse, and the innovation and economic and public health benefits that derive from it, rests on clear provenance and chain of evidence. Where AI tools are added as a component into the process of generating knowledge from knowledge, it should be possible for the AI to be placed into such a chain, necessitating appropriate and clear accountability mechanisms to achieve the same. Similarly, reproducibility is a quality that is key to the research enterprise and, if AI becomes part of that enterprise, it is critical that AI results are reproducible. An accountability framework with audits to validate reproducibility would be a useful instrument to ensure this.

Similarly, accountability tools should support and certify legal, licensed use of any copyrighted works used as inputs. STM's members are offering an ever-increasing number of products and services to make their content available, so that high-quality and accurate content sets can be used as training data in machine learning under various licensing schemes. The availability and accessibility of high-quality training data is vital for empowering AI developers with the licensed materials required to realize the benefits of AI to their full potential. Likewise, the wide array of licenses offered by publishers ensures that there are ample, accessible materials available for the continued training of both people and machines.

1d. Should AI audits or assessments be folded into other accountability mechanisms that focus on such goals as human rights, privacy protection, security, and diversity, equity, inclusion, and access? Are there benchmarks for these other accountability mechanisms that should inform AI accountability measures?

A broad goal for federal activity should be the prevention and combatting of misinformation, which undermines democratic processes, as well as the minimization of oversaturation of the information space, which results in wasted time and reduced productivity. Assessing and auditing AI for high-quality inputs, available under license, could support such goals.

1e. Can AI accountability practices have meaningful impact in the absence of legal standards and enforceable risk thresholds? What is the role for courts, legislatures, and rulemaking bodies?

At a minimum, clarity and transparency are required in the use of IP and copyright, and as part of any liability regime. AI systems can use huge volumes of copyright materials in the training process and as part of any commercial deployment, therefore transparency obligations will be necessary to enable rights holders to trace copyright infringements in content ingested by AI systems.

2. Is the value of certifications, audits, and assessments mostly to promote trust for external stakeholders or is it to change internal processes? How might the answer influence policy design?

Certifications and audits with respect to the inputs to AI systems and tools are required both to promote trust for external stakeholders and to ensure that internal processes are valid and rely on good inputs. A fundamental rule of any process, but particularly for algorithms and AI tools, is GIGO (Garbage In, Garbage Out). Responsible documentation of input and training data and information can help build trust for external stakeholders and support processes for the development of trustworthy and reliable AI tools.

3b. Accountability measures so that AI systems do not substantially contribute to harmful misinformation, disinformation, and other forms of distortion and content-related harms:

To prevent misinformation or disinformation, it is critical that quality, vetted inputs are used. As noted, these include those works produced and made available by scholarly publishers. We again strongly encourage the use of the “Version of Record”, which contains the latest peer-reviewed information and is preserved with integrity with any corrections, to promote accuracy and reduce the risk of erroneous outputs. Accountability that provides a trail of provenance, which ensures as well that any inputs have not been retracted or otherwise overturned, will help to minimize the threat of content-related harms.

In the context of scholarly publishing, STM recommends two things:

- An audit mechanism to validate that AIs operating on scientific content do not substantially alter their meaning and are able to provide a balanced summary of possibly different viewpoints in the scholarly literature. It should not be the role of an AI to say ‘who is right’ in an academic debate, but to provide information based on state-of-the-art research, based on licensed use of the Version of Record, and reflecting any corrections or retractions.
- An accounting with respect to provenance. This could be a simple list of references and inputs, or potentially something more sophisticated, to enable users or auditors to track back assertions generated by AIs to the original literature inputs.

3f. Accountability measures so that there are adequate human alternatives, consideration, and fallbacks in place throughout the AI system lifecycle:

Where AI tools are used in the domain of scholarly communications, there should always be the possibility for humans to review chains of evidence across the literature. Human intelligence is critical to discerning the validity of conclusions, so there should be no AI ‘black box’ in the chain of scholarly discovery.

5. Given the likely integration of generative AI tools such as large language models (e.g., ChatGPT) or other general-purpose AI or foundational models into downstream products, how can AI accountability mechanisms inform people about how such tools are operating and/or whether the tools comply with standards for trustworthy AI?

These tools are already confusing the public by reporting information that is inaccurate, inappropriate, or does not exist. These have been referred to as “hallucinations,” but in fact dilute trust in science and scholarly literature by sometimes even referring to or citing references that do not exist. At a minimum, AI tools and services should be required to accurately represent references and the sources of information that the tool has relied on, preferably under license. Where a tool or service explicitly relies on scholarly literature as input and generates text based on that, end-users should be informed about which AI tools have been utilized in the process and what the original corpus of input data has been.

15a. Where in the value chain should accountability efforts focus?

In the context of AI applied to scholarly literature, we reinforce the importance of describing the set of literature on which AI models have been trained and validated. In addition, for some tools, accountability should include reproducibility, that is, being able to reproduce a specific set of results from given sets of inputs, to enable audits and advance trust in the AI tool or service.

20. What sorts of records (e.g., logs, versions, model selection, data selection) and other documentation should developers and deployers of AI systems keep in order to support AI accountability?

As above, in the context of AI applied to the scholarly literature, accountability records should also include all source material that has been used – both in the application of an AI tool or system but also in its training and validation, as well as the licenses under which it has been provided. This should be maintained for the duration of the existence of the AI system.

22. How should the accountability process address data quality and data voids of different kinds? For example, in the context of automated employment decision tools, there may be no historical data available for assessing the performance of a newly deployed, custom-built tool. For a tool deployed by other firms, there may be data a vendor has access to, but the audited firm itself lacks. In some cases, the vendor itself may have intentionally limited its own data collection and access for privacy and security purposes. How should AI accountability requirements or practices deal with these data issues? What should be the roles of government, civil society, and academia in providing useful data sets (synthetic or otherwise) to fill gaps and create equitable access to data?

As noted, data quality can be addressed by articulating provenance and relying upon high-quality, validated inputs. In particular, when scholarly literature is used as an input, developers should use Versions of Record of articles made available under license for training and development of AI tools and systems. The Version of Record is the most thoroughly vetted version of a research publication, having been through all stages of the peer-review and publication process and subject to ongoing curation in case issues come to light post-publication (e.g., a corrigendum, erratum, or – in extreme cases – a retraction).

Similarly, transparency in accounting for the works used can help address data voids. There is a potential for developers, either by design or negligence, to “cherry pick” information that could be used as inputs for an AI and that would create bias or inaccurate information due to the missing sources of information. Publishers are happy to work with developers to ensure appropriate access under appropriate terms and licenses that would ensure the quality of output to fill any gaps that are identified.

27: What is the role of intellectual property rights, terms of service, contractual obligations, or other legal entitlements in fostering or impeding a robust AI accountability ecosystem? For example, do nondisclosure agreements or trade secret protections impede the assessment or audit of AI systems and processes? If so, what legal or policy developments are needed to ensure an effective accountability framework?

STM’s member publishers are at the forefront of digital innovation. They provide organized, quality information and content that is tagged and enriched with a variety of metadata. The accuracy of the scientific record supported by our publishers helps to ensure that machine learning and artificial intelligence applications have an appropriate depth and accuracy of information to consume. The availability and accessibility of high-quality training data is essential for empowering AI developers with the licensed materials required to achieve their goals in a responsible way. The wide array of licenses

that publishers offer ensures that both people and machines have options that make sense for their particular set of objectives.

To underline this point, the high-quality input STM's members provide, made possible by the incentive of copyright, is both an engine of free expression and a driver of innovation for enabling accountable AI as well as more broadly. We align with the CCC's response to this question, under which it notes that "[r]espect for copyright is consistent with and will foster a robust accountability ecosystem." We also agree with CCC that licensing enables end users to reuse high-quality content in scientific, technical, and medical manuscript Versions of Record. This is the only version of an article that includes all authoritative identifying metadata along with the final peer-reviewed, verified content. STM also aligns with the AAP's response to this question, which highlights that AI trained on anything other than the Version of Record "could create serious and cascading scientific or medical errors in AI generated outputs." Protecting and enforcing copyright bolsters the promotion of transparency, and in turn supports STM's mission, which is to advance trusted research worldwide.

In conclusion, STM and its members are keenly interested in accountability for AI and other AI policy and stand ready to engage with the federal government in support of trust and integrity in these technologies.

Respectfully submitted,



Dr. Caroline Sutton
CEO
STM

About STM

At STM we support our members in their mission to advance trusted research worldwide. Our more than 140 members collectively publish 66% of all journal articles and tens of thousands of monographs and reference works. As academic and professional publishers, learned societies, university presses, start-ups and established players, we work together to serve society by developing standards and technology to ensure research is of high quality, trustworthy and easy to access. We promote the contribution that publishers make to innovation, openness and the sharing of knowledge and embrace change to support the growth and sustainability of the research ecosystem. As a common good, we provide data and analysis for all involved in the global activity of research.

The majority of our members are small businesses and not-for-profit organizations, who represent tens of thousands of publishing employees, editors, reviewers, researchers, authors, readers, and other professionals across the United States and world who regularly contribute to the advancement of science, learning, culture and innovation throughout the nation. They comprise the bulk of a \$25 billion publishing industry that contributes significantly to the U.S. economy and enhances the U.S. balance of trade.