

AI *et al.*: Machines Are About to Change Scientific Publishing Forever

Cite This: *ACS Energy Lett.* 2023, 8, 878–880

Read Online

ACCESS |

Metrics & More

Article Recommendations

Supporting Information

Artificial intelligence (AI)-powered text generation will change scientific publishing fundamentally. In the past year, multiple AI systems have showcased production of visual and textual content increasingly indistinguishable from human-generated work, creating almost overnight new possibilities for intellectual workers, and at the same time raising similarly potent concerns. While artists and journalists are more evidently at the forefront of this incipient revolution, it is not hard to imagine a researcher looking away from the frustratingly sparse draft of a research article and wondering: “Could a machine write it for me?” (Figure 1).

This question might have passed for a flight of fancy until recently, as machine-generated scientific arguments were easily distinguishable from human output, and paper-generating software mainly highlighted the permeability of the peer-review process to nonsensical papers.¹ However, these technologies have progressed so rapidly that we have likely entered a new phase, one in which machine-generated text can be integrated in human-generated scientific articles in a seamless fashion. To help illustrate the point in concrete terms, let’s directly ask one of the involved parties:

Question: Can AI write a scientific paper? Answer in two paragraphs.

ChatGPT’s answer:

Yes, AI can write a scientific paper. In fact, there are already AI systems that can write scientific papers by using advanced natural language processing techniques to analyze existing scientific literature and generate text that is similar to human-written scientific papers. This allows AI to quickly and efficiently generate a large amount of text based on a given topic or research area, which can be useful for researchers who want to quickly generate a draft of a paper or for organizations that need to produce a large number of papers on a particular topic.

However, it is important to note that these AI-generated papers are typically not original research, but rather are based on existing research in a particular field. This means that they may not contain any new insights or findings, and are not likely to be published in academic journals. Additionally, the accuracy and validity of the information in the paper may not be as high as a human-written paper, so it is important for a human researcher to review and verify the information before it is published.

This well-structured answer makes reasonable points on both the potential and limitations of its author. As we write, the “assisted-driving” approach suggested by ChatGPT (AI-




Figure 1. Artificial Intelligence assisting the writing process, as imagined by the Midjourney text-to-image generator.

generated text with human supervision) seems to be employed by a growing number of researchers,² promising to free researchers’ time from the burden of scientific writing and get them back to the science. As a matter of fact, ChatGPT’s answer appears conservative, downplaying the role that AI could play in providing scientific insight. With their ability to take into account large swathes of scientific literature and condense them into concise chunks, AI language models could be used by researchers to familiarize themselves with new topics and double-check the completeness of their literature overviews. Since AI models have already made their way into scientific literature search engines,² it is easy to imagine their evolution

Received: December 13, 2022

Accepted: December 22, 2022

Published: January 4, 2023



into brainstorming tools, leveraging their extensive literature knowledge to produce suggestions based on recurrent scientific themes.

Additionally, while current AI language models do not produce statements reaching beyond the content of their training sets, we should ask ourselves whether this is a fundamental feature of the technology or just a temporary limitation. Recent breakthroughs in the use of AI in complex strategic games have highlighted the surprising ease with which AI can outcompete humans in problems considered intractable with computational approaches.^{3,4} Similar breakthroughs in the use of AI for scientific progress might come from the combination of (i) precise goals, i.e., a clear definition of what we consider a successful scientific observation, (ii) algorithms capable of efficiently optimizing its output for these goals, and (iii) structured and accessible scientific data. In this line of thought, we can imagine AI systems proposing new experiments and new descriptions of observed phenomena and arranging data in figures to support their conclusions. An AI system capable of producing original scientific work could revolutionize the whole scientific endeavor—for example, by being less tied than humans to the boundaries of scientific disciplines, bringing multidisciplinary science to new heights.

However, as we take the first steps down the slope of the AI revolution, it is worth envisioning the risks awaiting us at the bottom.

Probably the most immediate concern is that the sandwich paper⁵ and paper-mill systems could become supercharged when coupled with advanced AI language models. A text-generation system combining speed of implementation with eloquent and structured language could enable a leap forward for the serialized production of scientific-looking papers devoid of scientific content, increasing the throughput of paper factories and making detection of fake research more time-consuming. The proliferation of these practices has already had an impact on the publishing world despite the lack of sophisticated approaches,¹ and the problem will likely get worse if AI joins the fray. Scaling up efforts to automatically detect AI content might play an important role in tackling the issue.⁶ However, it is clear that these approaches would require the establishment of clear definitions for fraudulent and legitimate use of AI content. AI can already write articles that would, in the opinion of the authors, pass peer review for a Perspective article (see the [Supporting Information](#) for a paper on lead toxicity in perovskite devices, generated by AI). This ability makes the urgent need for a code of conduct for the use of AI-generated text in scientific literature abundantly clear.

On a more abstract level, there are questions regarding the originality of AI-science. Is an AI-generated scientific text original, despite being the product of training over human-made original work? Who is the intellectual owner of the content? These questions echo the ongoing debate surrounding AI-generated art, and suggest similar concerns on who stands to benefit from AI intellectual production. If an AI system is able to gather its information from the entire literature, will it be able to trace its conclusion back to the most relevant scientific work and acknowledge it? Or does the production of scientific content from delocalized literature sources push toward a new paradigm with a decreased emphasis on human authorship? Can we apply straightforwardly the current model of scientific publishing, based on individual contributions and intellectual property transfer from authors to journals, to an increasingly automated environment?

Increased AI-automation might also encroach on the originality of writing styles, supplanting the idiosyncrasies of human authors with a homogeneous style. Is the loss of individual stylistic traits, often culturally influenced, a step closer to mutual comprehension or the removal of a desired feature? This aspect might depend on the degree to which AI systems operate a “regression to the mean” of their training content, flattening potentially valuable deviations. Managing the amount of randomness utilized in the generation of content might offer ways to prevent the homogenization of AI production, but it is worth pondering whether these systems can replicate the diversity of voices and viewpoints provided by human scientists.

While current AI language models do not produce statements reaching beyond the content of their training sets, we should ask ourselves whether this is a fundamental feature of the technology or just a temporary limitation.

Finally, concerning the long-term consequences of the integration of AI in scientific publishing, it is worth considering the vulnerability of AI models to biases. Human researchers are far from being free of biases, but the presence of a plurality of voices helps in restricting the impact of individual biases, and traceability of authorship makes it possible to actively tackle biases in the system. Using a small number and diversity of AI systems for scientific production has the potential to amplify the scale of biases, propagating them efficiently on a global level and avoiding the mechanism of individual scrutiny often used to identify these biases. Since the behavior of these systems depends strongly on their training, it is crucial to understand the composition of the training sets and ensure they are constantly updated with all accepted forms of scientific production. And even if the AI models would be able to gather all the available literature, what would their approach toward consensus-defying reports be? Are they poised to stand on the side of commonly held position and silence singular points of view? Finally, would an AI-centric system of scientific content production be vulnerable to deliberate biasing by malicious actors? These prospects suggest that it is crucial to consider and potentially regulate the openness of the AI models and training sets used, and to start worrying about the emergence of malicious strategies to amplify the relevance of selected opinions.

Some of these concerns might be mitigated with small adjustments, others might require systemic changes, and others still might prompt us to restrict the use of these new technologies. What appears clear is that AI will make its way into the practice of scientific publishing soon, if it has not already done so (how would we know?). In doing so it will change science appreciably and permanently. When computers made their appearance, they entirely changed the way science is performed and communicated—such changes required adjustment and adaptation. The appearance of advanced AI promises to ratchet up the degree of change and, at the same time, the amount of adaptation required. Like our forest-dwelling ancestors that discovered fire, we need to be mindful of the unwanted consequences of our exciting advances, to reap their benefits without setting our homes ablaze.

Gianluca Grimaldi orcid.org/0000-0002-2626-9118

Bruno Ehrler orcid.org/0000-0002-5307-3241

■ ASSOCIATED CONTENT

SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acsenerylett.2c02828>.

Example of a Perspective article produced by prompting GPT-3 on the topic of lead toxicity in perovskite photovoltaics, followed by an Appendix listing the prompts used and the minor edits that the authors implemented on GPT-3's output ([PDF](#))

■ AUTHOR INFORMATION

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acsenerylett.2c02828>

Notes

The authors generated the paragraphs labeled as “ChatGPT’s Answer” in the Viewpoint and the sample article in the Supporting Information (SI) in part with GPT-3, OpenAI’s large-scale language-generation model. The AI-generated text included in the Viewpoint was left unaltered. Upon generating draft language for the SI article, the authors reviewed, edited, and revised the language to their own liking, reporting every change from the original text in the Appendix section. The authors do not take responsibility for the content of the AI-written text, which is used merely to illustrate the level of sophistication of the language-generation model.

The authors declare no competing financial interest.

Views expressed in this Viewpoint are those of the authors and not necessarily the views of the ACS.

■ ACKNOWLEDGMENTS

We thank Prashant Kamat, Imme Schuringa, Jeroen de Boer, and Moritz Schimdt for the stimulating discussions on the recent developments of AI-generation of scientific content and on the long-term impact of AI on scientific writing.

■ REFERENCES

- (1) Van Noorden, R. Hundreds of Gibberish Papers Still Lurk in the Scientific Literature. *Nature* **2021**, *594* (7862), 160–161.
- (2) Hutson, M. Could AI Help You to Write Your next Paper. *Nature* **2022**, *611* (7934), 192–193.
- (3) Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; Dieleman, S.; Grewe, D.; Nham, J.; Kalchbrenner, N.; Sutskever, I.; Lillicrap, T.; Leach, M.; Kavukcuoglu, K.; Graepel, T.; Hassabis, D. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature* **2016**, *529* (7587), 484–489.
- (4) Meta Fundamental AI Research Diplomacy Team (FAIR); Bakhtin, A.; Brown, N.; Dinan, E.; Farina, G.; Flaherty, C.; Fried, D.; Goff, A.; Gray, J.; Hu, H.; Jacob, A. P.; Komeili, M.; Konath, K.; Kwon, M.; Lerer, A.; Lewis, M.; Miller, A. H.; Mitts, S.; Renduchintala, A.; Roller, S.; Rowe, D.; Shi, W.; Spisak, J.; Wei, A.; Wu, D.; Zhang, H.; Zijlstra, M. Human-Level Play in the Game of Diplomacy by Combining Language Models with Strategic Reasoning. *Science* **2022**, *378* (6624), 1067–1074.
- (5) Jin, S. Fewer Sandwich Papers, Please. *ACS Energy Lett.* **2022**, *7* (10), 3727–3728.
- (6) Jawahar, G.; Abdul-Mageed, M.; Lakshmanan, L. V. S. Automatic Detection of Machine Generated Text: A Critical Survey. *arXiv*, Nov 2, 2020, arXiv:2011.01314 [cs.CL].