

Artificial Intelligence in Scientific Publishing: A Tool to Be Governed, Not a Substitute for Scientific Judgment

Claudio Colosio

Post Graduate School in Occupational Medicine, University of Milano, 20142 Milan, Italy; claudio.colosio@unimi.it
 Received: 14 August 2026; Accepted: 21 September 2026; Published: 15 October 2026

How To Cite: Colosio, C. Artificial Intelligence in Scientific Publishing: A Tool to Be Governed, Not a Substitute for Scientific Judgment. *Work and Health* **2026**, 2(4), 23. <https://doi.org/10.53941/wah.2026.100023>

Technological innovation has rarely entered human society without resistance. Some of the concerns that accompanied technologies now regarded as indispensable may appear almost incomprehensible today, although they were taken very seriously at the time.

The Industrial Revolution provides perhaps the best-known example. In early nineteenth-century England, the Luddites destroyed textile machinery that they perceived as threatening their jobs, skills, and livelihoods. Their reaction was not simply an irrational fear of machines: it also reflected legitimate concerns about profound changes in working conditions and the distribution of the benefits of technological progress. Nevertheless, mechanization continued and eventually transformed manufacturing and society.

Other innovations followed similar paths. Railways were initially regarded by some as potentially dangerous to human health. The automobile encountered resistance and restrictive legislation: the British Locomotive Act of 1865 famously required, under certain circumstances, a person carrying a red flag to precede a self-propelled vehicle. Electricity, elevators, vaccination and, much more recently, the Internet all encountered scepticism, resistance or fear before becoming part of everyday life.

History does not teach us that every technological innovation is necessarily beneficial. There have been technologies whose harmful consequences became apparent only after widespread adoption. Nor should legitimate concerns about technological change be dismissed as another form of Luddism. History does, however, suggest that genuinely useful technologies tend eventually to become incorporated into ordinary life and professional practice.

Artificial intelligence (AI), and particularly generative AI based on large language models (LLMs), appears likely to follow this pattern.

The scientific community is now debating whether and how AI should be used in research, scientific writing, peer review and editorial decision-making. Positions range from enthusiasm to considerable concern. Both attitudes are understandable. AI systems can perform tasks that only a few years ago would have seemed extraordinary. At the same time, their limitations—including inaccurate information, fabricated references, bias, lack of transparency and the generation of extremely convincing but incorrect statements—are equally real [1–4].

The question, therefore, should probably no longer be whether scientists will use AI. They already do, and its use will almost certainly increase.

A more useful question is: what should we allow AI to do for us, and what must we never delegate to it?

1. Assistance and Delegation Are Not the Same Thing

This distinction, in our view, is fundamental.

AI can be an extraordinarily useful scientific assistant.

Consider the preparation of a manuscript. AI can identify grammatical and spelling errors, improve the readability of English, suggest more concise formulations, identify repetitions, detect inconsistencies between different sections of a manuscript and help organize complex material. It can assist in searching the literature and in identifying potentially relevant references, provided that these references are subsequently checked against the original sources.



For scientists whose first language is not English, this may represent a particularly important development. Scientific merit and proficiency in English are obviously not the same thing, although the two have sometimes been implicitly confused in international scientific publishing. An excellent researcher may have an important scientific message and nevertheless encounter difficulties expressing it in polished academic English. Properly used, AI can reduce this linguistic disadvantage.

There is an interesting connection here with the mission of an international journal such as *Work and Health*. Occupational health is intrinsically global, but scientific resources and opportunities to publish internationally are not equally distributed. The possibility of using AI to overcome at least part of the linguistic and technical barriers to scientific communication should therefore not be underestimated.

2. AI as a Potential Equalizer in Scientific Publishing

The potential contribution of AI may be particularly relevant for researchers working in low- and middle-income countries. Scientific manuscripts originating from these settings may contain valuable data and address occupational and environmental health problems of considerable local and global importance, while their chances of publication may nevertheless be reduced by barriers that are only partly related to the scientific value of the research.

Language is an obvious example. Limited proficiency in scientific English can obscure otherwise valid scientific content, make a manuscript difficult to evaluate, and ultimately contribute to rejection. Generative AI can substantially reduce this disadvantage by improving grammar, syntax and clarity without necessarily altering the underlying scientific message.

A second, and perhaps equally important, example concerns statistical analysis. Researchers working in institutions with limited methodological support may have difficulty identifying the most appropriate statistical approach, even when the underlying principles are well established. The choice of a statistical test depends on features such as study design, type and distribution of variables, independence or dependence of observations, sample size, assumptions of the model, and the hypothesis being tested. These principles are described in statistical textbooks and methodological guidelines—the established rules of the discipline—but access to experienced biostatistical support remains very uneven across institutions and countries.

Properly interrogated, AI can help researchers navigate these established principles. It can suggest potentially appropriate statistical approaches, explain why one test may be preferable to another, identify assumptions that need to be checked, and even draft preliminary code for statistical environments such as R or suggest procedures to be implemented in packages such as SPSS. The proposed analysis can then be performed, reproduced and critically verified using validated statistical software and, whenever appropriate, reviewed by a statistician.

Here again, assistance must not become delegation. AI should not be asked to search for an analysis that produces the desired result, nor should its statistical recommendations be accepted without understanding and verification. The researcher must understand why a particular method has been selected and remains responsible for its appropriateness and interpretation.

Used within these boundaries, AI could become more than a convenient writing tool. It could contribute to the democratization of scientific publishing by reducing some of the linguistic and technical inequalities that limit participation in international research. This may be especially relevant to occupational and environmental health, where important evidence often originates precisely from countries with limited research resources.

At the same time, editors and reviewers may need to reconsider the weight traditionally given to the formal quality of scientific English. In an era in which AI can produce polished prose, linguistic elegance will increasingly be neither evidence of scientific quality nor, by itself, evidence of inappropriate AI use. The underlying science must remain the primary object of evaluation.

The same reasoning can apply to peer review.

An experienced reviewer may use AI to help organize comments, identify passages that deserve closer attention, improve the language of a review, or question whether some aspect of the manuscript has been overlooked. Used in this manner—and only where journal policies and confidentiality requirements permit it—AI may act almost as an additional interlocutor with whom the reviewer tests his or her reasoning.

But this is where a clear boundary must be drawn.

Assistance is not delegation.

The decision as to whether a scientific hypothesis is meaningful cannot be delegated to AI. The interpretation of epidemiological findings cannot be delegated to AI. The assessment of biological plausibility, the distinction between association and causation, the evaluation of the relevance of an exposure, and the judgment of whether the conclusions of a study are actually supported by its results cannot be delegated to AI.

And, ultimately, the decision as to whether a manuscript deserves publication cannot be delegated to AI.

These are strategic scientific decisions. They require knowledge, experience, critical thinking, awareness of uncertainty and, not infrequently, the intuition that develops after decades of scientific work.

AI may contribute information to these decisions.

It must not make them for us.

3. AI, Scientific Creativity and the Risk of Intellectual Convergence

There is another, less immediately visible risk in delegating scientific reasoning to AI. Generative AI operates primarily on knowledge that has already been produced. It is exceptionally powerful at identifying patterns, summarizing prevailing interpretations, connecting existing information and suggesting plausible extensions of what is already known. But science does not progress only by organizing existing knowledge. It also progresses by questioning it.

Many important scientific advances have originated from observations that did not fit the prevailing explanation, from questions that had not previously been asked, or from hypotheses that initially appeared improbable in the light of established knowledge. An AI system may certainly generate apparently novel hypotheses by combining existing concepts in new ways. What it cannot replace is the scientist's capacity to recognize that an unexpected observation matters, to challenge an accepted paradigm, or to decide that an apparently marginal finding deserves to become the starting point of an entirely new line of investigation.

This distinction is particularly important when AI is used for literature searching, evidence synthesis and peer review. If these activities are delegated rather than assisted, there is a risk of intellectual convergence: what is already prominent, frequently cited and consistent with prevailing interpretations may become still more visible, while unconventional hypotheses, negative findings, less represented scientific traditions and evidence lying outside the expected conceptual framework may receive less attention. In this way, an instrument designed to expand access to knowledge could paradoxically narrow the range of ideas that scientists actively consider.

The danger is therefore not only that AI may provide an incorrect answer. A subtler danger is that it may repeatedly provide the most plausible answer within the boundaries of existing knowledge. Scientific creativity sometimes begins precisely where plausibility ends: with the question that has not yet been asked, the observation that does not fit, or the hypothesis that initially appears unlikely.

AI can help us explore the map of existing knowledge. Scientists must remain responsible for deciding when it is necessary to leave the map.

4. A Remarkably Competent Assistant that Can also Be Wrong

This distinction becomes even more important because of one characteristic of generative AI: it can be extraordinarily persuasive.

A traditional computer error often looks like an error. An AI error may look like an authoritative scientific statement.

LLMs can produce beautifully structured paragraphs, convincing explanations and plausible references that may nevertheless be partly or entirely incorrect. The literature has extensively discussed this problem, usually under the perhaps anthropomorphic term "hallucination" [3,4].

This creates an unusual risk in science. Poorly written nonsense is relatively easy to identify. Eloquently written nonsense is much more dangerous.

The solution, however, is not necessarily to reject the tool. It is to understand what the tool can and cannot do.

We already accept this principle for other technologies used routinely in research. Statistical software performs calculations that few researchers would now perform manually. But the software cannot decide whether the chosen statistical model appropriately addresses the scientific question. Reference-management software can generate a bibliography in seconds, but it cannot guarantee that the cited paper supports the statement made in the manuscript. A spreadsheet can process thousands of observations, but responsibility for the validity of the data remains with the researcher.

AI should be approached in the same spirit, although its apparent intellectual autonomy makes critical supervision even more important.

The principle could be summarized simply: AI-generated output should be considered a proposal, not evidence.

Every important factual statement should be verified. Every reference should be checked against the original publication. Every scientific interpretation should be critically assessed. And every conclusion should ultimately be the conclusion of the human author, reviewer or editor.

5. Peer Review Presents a Particular Challenge

The use of AI in peer review deserves special consideration. Reviewing a scientific manuscript is not simply an exercise in identifying errors. A good reviewer tries to understand what the authors attempted to demonstrate, whether their methods allow them to do so, whether alternative explanations have been adequately considered and whether the results meaningfully advance existing knowledge. This requires judgment.

Confidentiality is a separate and fundamental issue. A manuscript submitted for peer review is privileged, unpublished information, provided to reviewers for a specific purpose and under an implicit or explicit obligation not to disclose its contents. Current editorial recommendations therefore warn reviewers against uploading submitted manuscripts to AI systems when confidentiality cannot be assured [5,6].

At first sight, however, this deserves a more precise question. If the material that I provide to an AI system truly remains between me and the tool, in the same way that a document processed by other software remains within a protected environment, why should confidentiality necessarily have been breached? The answer is that the problem is not the mere fact that AI has been used. The problem is whether the reviewer can actually be assured of what happens to the information after it has been entered into the system.

With a conventional local tool, the boundaries may be relatively clear. With an external generative AI service, they may not be. Depending on the product, account, contractual terms, institutional arrangements and settings, prompts or uploaded material may be transmitted to external servers, retained for some period, logged, processed by third parties, or handled under conditions that the reviewer has not independently established. Conversely, an AI system operating in a publisher-controlled or institutionally approved environment may provide contractual and technical safeguards that make confidentiality reasonably assured. Thus, the relevant distinction is not simply between ‘AI’ and ‘no AI’, but between environments in which confidentiality is demonstrably protected and those in which it is not.

This is, importantly, also the logic of the current ICMJE formulation: the prohibition concerns uploading manuscripts into AI systems where confidentiality cannot be assured. It is therefore a risk-based principle rather than a claim that every use of AI inherently constitutes disclosure [5]. Nevertheless, reviewers remain bound by the journal’s policy and should obtain permission where that policy requires it, because the manuscript is entrusted to the reviewer by the journal and the authors, not owned by the reviewer.

The expression ‘use of AI in peer review’ also encompasses very different activities. Uploading a complete manuscript and asking an external system to produce a review is very different from asking a general methodological question without revealing manuscript content, or from using AI to improve the English and organization of comments written by the reviewer. These activities do not carry the same confidentiality risk.

The relevant questions are therefore not simply ‘Was AI used?’ but: What information was provided to the AI system? Where and how was it processed? Can confidentiality be reasonably assured? Was its use consistent with the journal’s policy? And for what purpose was the AI used?

This more nuanced approach is likely to become increasingly important as AI systems evolve. A prudent principle is that unpublished manuscripts, data, figures and other confidential information should not be transferred to an AI environment unless its confidentiality safeguards are known to be adequate and its use is permitted by the journal. Where these conditions are satisfied, the confidentiality objection is substantially different; where they are not, the reviewer should not upload the material.

Once again, governance rather than prohibition provides the more useful framework. AI may support peer review, but confidentiality must be demonstrably protected and the reviewer’s responsibility cannot be transferred to the tool.

6. From Prohibition to Governance

Faced with these problems, prohibition may initially appear to be the safest solution. We believe that, in the long term, it is unlikely to be either realistic or desirable.

Science has repeatedly incorporated technologies that changed the way researchers work. Electronic databases transformed literature searches. Statistical packages changed data analysis. Reference managers changed manuscript preparation. Electronic submission systems changed publishing. Automatic spelling and grammar correction became so normal that almost nobody considers their use ethically problematic.

Generative AI is clearly more powerful and raises more complex ethical questions than these earlier tools. Nevertheless, attempting simply to exclude it from scientific activity would probably ignore both its potential benefits and the reality of its rapidly increasing use.

The appropriate response should therefore be governance rather than prohibition.

Three principles appear particularly important.

The first is transparency.

Substantial use of generative AI in preparing scientific work should be disclosed. The reader and the editor should be able to understand which tool was used and for what purpose. Transparency should not be regarded as an admission of improper behaviour. On the contrary, transparent disclosure distinguishes responsible use from concealed delegation.

The second principle is verification.

AI output must never acquire scientific authority merely because it is well expressed. Facts must be checked. References must be verified. Interpretations must be challenged. Researchers should remain particularly alert to the possibility that AI systems reproduce biases contained in their training material or generate apparently reasonable answers when the available evidence is uncertain.

The third, and most important, principle is human accountability.

An AI system cannot be accountable for a scientific publication. It cannot defend a methodological decision before the scientific community. It cannot accept responsibility for an incorrect conclusion. It cannot respond to an allegation of scientific misconduct. For these reasons, major international recommendations agree that AI cannot qualify as an author [5,6].

The responsibility for a manuscript belongs to its authors.

The responsibility for a review belongs to the reviewer.

The responsibility for an editorial decision belongs to the editor.

This chain of responsibility should never be interrupted.

7. Keeping the Scientist in Command

The debate surrounding AI is sometimes presented as a choice between embracing the technology enthusiastically or defending traditional science against it.

We believe this is a false choice.

AI is a tool. A uniquely powerful tool, certainly, but still a tool.

Its scientific value will depend to a considerable extent on the person using it. A researcher with limited knowledge who delegates critical thinking to AI may produce an elegant but scientifically weak manuscript. An experienced scientist who interrogates AI critically, checks its answers, rejects inappropriate suggestions and uses it to accelerate non-strategic tasks may substantially improve efficiency without surrendering intellectual control.

The same applies to reviewers and editors.

Indeed, the greatest risk may not be that AI becomes too intelligent, but that scientists become intellectually passive in using it.

The purpose of AI should be to expand our ability to think, not to relieve us of the obligation to think.

This distinction is particularly relevant in occupational and environmental health. Our discipline routinely deals with incomplete evidence, complex exposure scenarios, multiple causal pathways, uncertain risk estimates and decisions that may have substantial consequences for workers, employers and society. Algorithms may help us organize and analyse information, but scientific judgment remains indispensable.

Antonio Mutti recently addressed this issue in *La Medicina del Lavoro*, emphasizing the need to preserve disciplinary integrity and editorial standards as AI enters occupational medicine [7]. His position reflects a broader evolution in scientific publishing: the discussion is gradually moving away from the simple question of whether AI should be allowed and toward the more difficult question of how its use should be governed.

The ICMJE recommendations updated in 2026 similarly identify validity, confidentiality, transparency and accountability as central principles governing AI use in scientific publishing [5].

These principles offer a sensible path forward.

AI should not replace the scientist.

It should make a good scientist more effective.

Perhaps, therefore, the rule for the coming years can be expressed in one sentence:

Use AI to help you think, but never delegate to it the responsibility of thinking for you.

The history of technological progress suggests that useful innovations eventually become part of ordinary professional practice. Artificial intelligence will probably be no exception. The challenge for science is neither to resist it indiscriminately nor to surrender to it uncritically.

It is to learn how to govern it.

And, above all, to remain in command.

Conflicts of Interest

The author declares no conflict of interest.

Use of AI and AI-Assisted Technologies

ChatGPT (OpenAI) was used as an AI-assisted tool during the preparation of this Editorial to support literature identification and organization, manuscript structuring, critical discussion of alternative formulations, and English-language editing. Its use was limited to support functions and did not involve delegation of scientific or editorial judgment. All references and factual statements have been independently verified against the original sources before submission. The arguments, interpretation of the literature, scientific judgments, strategic decisions, and final conclusions remain entirely those of the author, who assumes full responsibility for the content of the Editorial.

References

1. Mutti, A. Hey James, Write an Editorial for “La Medicina del Lavoro”. *La Med. del Lav.* **2023**, *114*, e2023014. <https://doi.org/10.23749/mdl.v114i2.14451>.
2. Thirunavukarasu, A.J.; Ting, D.S.J.; Elangovan, K.; et al. Large language models in medicine. *Nat. Med.* **2023**, *29*, 1930–1940. <https://doi.org/10.1038/s41591-023-02448-8>.
3. Salvagno, M.; Taccone, F.S.; Gerli, A.G. Can artificial intelligence help for scientific writing? *Crit. Care* **2023**, *27*, 75. <https://doi.org/10.1186/s13054-023-04380-2>.
4. Alkaissi, H.; McFarlane, S.I. Artificial Hallucinations in ChatGPT: Implications in Scientific Writing. *Cureus* **2023**, *15*, e35179. <https://doi.org/10.7759/cureus.35179>.
5. International Committee of Medical Journal Editors. 의학학술지에 게재되는 학술 저작물의 생산, 보고, 편집 및 출판에 대한 권고안 (2024 년 1 월 개정): 국제의학학술지편집인위원회 [Recommendations for the conduct, reporting, editing and publication of scholarly work in medical journals (revised in January 2024): A Korean translation]. *Ewha. Med. J.* **2024**, *47*, e48. <https://doi.org/10.12771/emj.2024.e48>.
6. Zielinski, C.; Winker, M.A.; Aggarwal, R.; et al. Chatbots, generative AI, and scholarly manuscripts: WAME recommendations on chatbots and generative artificial intelligence in relation to scholarly publications. *Curr. Med. Res. Opin.* **2024**, *40*, 11–13. <https://doi.org/10.1080/03007995.2023.2286102>.
7. Mutti, A. Artificial Intelligence in Occupational Medicine: Upholding Disciplinary Integrity and Editorial Standards. *La Med. del Lav.* **2026**, *117*, 18996. <https://doi.org/10.23749/mdl.2026.18996>.