

# AI is transforming peer review — and many scientists are worried

Artificial intelligence software is increasingly involved in reviewing papers, provoking interest and unease.

26 March 2025

*Article written by Miryam Naddaf  
published by Nature  
News Features, 26 March 2025*




Illustration: Ibrahim Rayintakath

**Artificial intelligence software is increasingly involved in reviewing papers, provoking interest and unease.**

This February, ecologist Timothée Poisot was surprised when he read through the peer reviews of a manuscript he had submitted for publication. One of the referee reports seemed to have been written with, or perhaps entirely by, artificial intelligence (AI). It contained the telltale sentence, "Here is a revised version of your review with improved clarity

and structure", a strong indication that the text was generated by large language models (LLMs).

Poisot hasn't yet told the journal editor of his suspicions; he asked that the journal involved — which bans the use of LLMs in peer reviews — not be revealed in this article.

But in a [blogpost about the incident](#), he argued strongly against automated peer review. "I submit a manuscript for review in the hope of getting comments from my peers. If this assumption is not met, the entire social contract of peer review is gone," wrote Poisot, who works at the University of Montreal in Canada.

AI systems are already transforming peer review — sometimes with publishers' encouragement, and at other times in violation of their rules. Publishers and researchers alike are testing out AI products to flag errors in the text, data, code and references of manuscripts, to guide reviewers toward more-constructive feedback, and to polish their prose. Some new websites even offer entire AI-created reviews with one click.

But with these innovations come concerns. Although today's AI products are cast in the role of assistants, AI might eventually come to dominate the peer-review process, with the human reviewer's role reduced or cut out altogether. Some enthusiasts see the automation of peer review as an inevitability — but many researchers, such as Poisot, as well as journal publishers, view it as a disaster.

## My other editor is AI

Even before the appearance of ChatGPT and other AI tools based on LLMs, publishers had been using a variety of AI applications to ease the peer-review process for more than half a decade — including for tasks such as checking statistics, summarizing findings and easing the selection of peer reviewers. But the advent of LLMs, which mimic fluent

human writing, has changed the game.

In a [survey of nearly 5,000 researchers](#), some 19% said they had already tried using LLMs to 'increase the speed and ease' of their review. But the survey, by publisher Wiley, headquartered in Hoboken, New Jersey, didn't interrogate the balance between using LLMs to touch up prose, and relying on the AI to generate the review.



[Three AI-powered steps to faster, smarter peer review](#)

One study<sup>1</sup> of peer-review reports for papers submitted to AI conferences in 2023 and 2024 found that between 7% and 17% of these reports contained signs that they had been '[substantially modified](#)' by LLMs — meaning changes beyond spell-checking or minor updates to the text.

Many funders and publishers currently forbid reviewers of grants or papers from using AI, citing concerns about leaking confidential information if researchers load material into chatbot websites. But if researchers host offline LLMs on their own computers, then data aren't fed back into the cloud, says Sebastian Porsdam Mann at the University of Copenhagen, who studies the practicalities and ethics of using generative AI in research.

Using offline LLMs to rephrase one's notes can speed up and sharpen the process of writing reviews, so long as the LLMs don't "crank out a full review on your behalf", wrote Dritjon Gruda, an organizational-behaviour researcher at the Catholic University of Portugal in Lisbon, in a [Nature careers column](#).

But "taking superficial notes and having an LLM synthesize them falls far, far short of writing an adequate peer review", counters Carl Bergstrom, an evolutionary biologist at the University of Washington in Seattle. If reviewers start relying on AI so that they can skip most of the process of

writing reviews, they risk providing shallow analysis. “Writing is thinking,” Bergstrom says.

LLMs can certainly improve some reviewers’ style, says Porsdam Mann: this is unsurprising, given that some peer reviews are slapdash or poorly written. However, LLM output almost always contains errors, because the tools work by producing text that seems statistically likely on the basis of their training data and inputs — although researchers are finding ways to dampen error rates.

In many cases, the gap between humans and LLMs isn’t so great, according to a study that provided more than 300 US computational biologists and AI researchers with reviews of their own papers — some produced by human reviewers and others by GPT-4, one of the leading LLMs at the time<sup>2</sup>. Some 40% of respondents said the AI was either more helpful than the human reviews, or as helpful; and a further 42% that the AI was less helpful than many, but more helpful than some (see ‘Comparing AI and human peer review’).

Source: Ref. 2

## AI that goes beyond editing

The team behind the study that compared AI and human reviews, led by James Zou, a computational biologist at Stanford University, California, is now developing a reviewer ‘feedback agent’. It evaluates human review reports against a checklist of common issues — such as vague or inappropriate feedback — and, in turn, suggests how reviewers can improve their comments.

At a publisher innovation fair in London last December, many AI developers lined up to pitch products to improve peer review that do more than mere editing. One tool, called Eliza, launched last year by the firm World Brain Scholar (WBS) in Amsterdam, the Netherlands, makes

suggestions to improve reviewer feedback, recommends relevant references and translates reviews written in other languages into English. The tool is not meant to replace human peer reviewers, says WBS founder Zeger Karssen. "The tool will just analyse what the peer reviewer has written down," he says.

A similar tool is [Review Assistant](#), developed by multinational publishing-services firm Enago and Charlesworth. Initially, the tool used an LLM system to answer structured queries about a manuscript, which reviewers could then check or verify. But after talking to publishers, developers added a 'human first' mode, in which reviewers answer the queries and then have an AI tool look at their answers. The tool can "support reviewers to do what they may already be doing illegitimately, in a legitimate way", says co-developer Mary Miskin, global operations director at Charlesworth, who is based in Huddersfield, UK.

Another AI approach aims to free reviewers from the laborious parts of peer review. A start-up firm called Grounded AI, in Stevenage, UK, has developed a tool called [Veracity](#), which checks whether cited papers in manuscripts exist, and then — using an LLM — analyses whether the cited work corresponds to the author's claims. It functions like "the workflow that a motivated, rigorous human fact checker would go through if they had all the time in the world", says co-founder Nick Morley.



[What are the best AI tools for research? Nature's guide](#)

And a host of efforts have sprung up to apply LLM-assisted tools to existing papers — from software to spot image duplications, to statistics-checking programs. But researchers have expressed concerns that LLMs can be unreliable and that some apparent errors [could be false positives](#).

One AI review tool that's already in trials with publishers is [Alchemist](#)



Review, developed by Grounded AI and a company called Hum in Charlottesville, Virginia. The software's creators say that it can summarize core findings and methods and assess the novelty of research, as well as validate citations. They also say that reviewers can use the tool in a secure environment that protects the confidentiality of manuscripts and authors' intellectual property.

AIP Publishing, the publishing arm of the American Institute of Physics, headquartered in Melville, New York, is piloting a version of this software in two journals, says chief transformation officer Ann Michael. Journal editors will test a prototype of the tool and, at their discretion, allow some peer reviewers to try it. However, the publisher will not test the tool's ability to judge novelty, because internal surveys suggested editors didn't rate that as being as helpful as other features, Michael says. "We're trying to learn how to responsibly apply AI to peer review," she says, emphasizing that the tool is being used before human review, not to replace it.

Other publishers also told *Nature* that they were exploring developing in-house AI tools for peer review, but didn't say exactly what they were working on. Wiley, for instance, is "looking into various potential use cases for AI to strengthen peer review, including at the editor and reviewer levels", a spokesperson said.

A December 2024 study of guidelines at top medical journals<sup>3</sup> found that among large publishers, Elsevier currently bans reviewers from using generative AI or AI-assisted review, whereas Wiley and Springer Nature permit "limited use". Both Springer Nature and Wiley require the disclosure of any use of AI to support review, and forbid online uploading of manuscripts. (*Nature's* news team is editorially independent of its publisher.) The study noted that 59% of the 78 top medical journals with guidance on the matter ban AI use in peer review. The rest allow it, with varying requirements.

## AI-led review?

The most radical applications of AI in peer reviewing are tools that directly provide automated reviews of manuscripts. One example is [Paper-Wizard](#), which generates full-fledged multi-page reviews when a paper is uploaded, and checks for detailed aspects of methodological designs, such as statistical rigour. Its co-creator, cognitive neuroscientist Shane Ehrhardt in Brisbane, Australia, says that it is a 'pre-peer-review' product, intended to help authors with their own work.