Academics say this is just the beginning. The rising expectations of an emboldened labour movement were on full display on 23 December, when more than 35% of the members of two unions representing UC graduate students voted against accepting university officials' offer and ending the strike.

Missed opportunity?

One of the organizers of the vote-no campaign was Dylan Kupsh, a graduate researcher in computer science at UCLA. Kupsh was in close contact with union organizers at Columbia University, where student workers rejected an initial contract proposal and went on to secure further concessions after a ten-week strike that ended last January.

In the end, UC graduate students received a range of new benefits, including increased childcare subsidies; protections against bullying, discrimination and harassment; and a new schedule for salaries. Incoming graduate students, for example, will see their annual salary increase from around US$22,000 to $30,500. "We could have won a lot more, and it's sad that we didn't get there," Kupsh says. "We're going to have to repeat in another 2.5 years."

For Barry Eidlin, a sociologist at McGill University in Montreal, Canada, who studies the labour movement, the scale of the vote-no campaign is yet another sign of changing expectations in academia. "In the past, academic workers have felt like they should just keep their heads down and be grateful they have a job," he says. "The idea that people now expect more, and are willing to fight for more, seems to me a welcome shift in perspective."

ABSTRACTS WRITTEN BY CHATGPT FOOL SCIENTISTS

Researchers cannot always differentiate between AI-generated and original abstracts.

By Holly Else

An artificial-intelligence (AI) chatbot can write such convincing fake research-paper abstracts that scientists are often unable to spot them, according to a preprint posted on the bioRxiv server in late December[1]. Researchers are divided over the implications for science.

"I am very worried," says Sandra Wachter, who studies technology and regulation at the University of Oxford, UK, and was not involved in the research. "If we're now in a situation where the experts are not able to determine what's true or not, we lose the middleman that we desperately need to guide us through complicated topics," she adds.

The chatbot, ChatGPT, creates realistic text in response to user prompts. It is a 'large language model', a system based on neural networks that learn to perform a task by digesting huge amounts of existing human-generated text. Software company OpenAI, based in San Francisco, California, released the tool on 30 November, and it is free to use.

Since its release, researchers have been grappling with the ethical issues surrounding its use, because much of the chatbot's output can be difficult to distinguish from human-written text. Scientists have published a preprint[2] and an editorial[3] written by ChatGPT. Now, a group led by Catherine Gao at Northwestern University in Chicago, Illinois, has used ChatGPT to generate artificial research-paper abstracts to test whether scientists can spot them.

The researchers asked the chatbot to write 50 medical-research abstracts based on a selection published in *JAMA*, *The New England Journal of Medicine*, *The BMJ*, *The Lancet* and *Nature Medicine*. They then compared these with the

> "If the experts are not able to determine what's true, we lose the middleman to guide us through complicated topics."

original abstracts by running them through a plagiarism detector and an AI-output detector, and asked a group of medical researchers to spot the fabricated abstracts.

Under the radar

The ChatGPT-generated abstracts sailed through the plagiarism checker: the median originality score was 100%, which indicates that no plagiarism was detected. The AI-output detector spotted 66% of the generated abstracts. But the human reviewers didn't do much better: they correctly identified only 68% of the generated abstracts and 86% of the genuine abstracts. They incorrectly identified 32% of the generated abstracts as being real and 14% of the genuine abstracts as being generated.

Wachter says that, if scientists can't determine whether research is true, there could be "dire consequences". As well as being problematic for researchers, who could be pulled down flawed routes of investigation, because the research they are reading has been fabricated, there are "implications for society at large because scientific research plays such a huge role in our society". For example, it could mean that research-informed policy decisions are incorrect, she adds.

But Arvind Narayanan, a computer scientist at Princeton University in New Jersey, says: "It is unlikely that any serious scientist will use ChatGPT to generate abstracts." He adds that whether generated abstracts can be detected is "irrelevant". "The question is whether the tool can generate an abstract that is accurate and compelling. It can't, and so the upside of using ChatGPT is minuscule, and the downside is significant," he says.

Irene Solaiman, who researches the social impact of AI at Hugging Face, an AI company with headquarters in New York and Paris, has fears about any reliance on large language models for scientific thinking. "These models are trained on past information and social and scientific progress can often come from thinking, or being open to thinking, differently from the past," she adds.

The authors suggest that those evaluating scientific communications, such as research papers and conference proceedings, should put policies in place to stamp out the use of AI-generated texts. If institutions choose to allow use of the technology in certain cases, they should establish clear rules around disclosure. This month, the Fortieth International Conference on Machine Learning — which will be held in Honolulu, Hawaii, in July — announced that it has banned papers written by ChatGPT and other AI language tools.

Solaiman adds that in fields where fake information can endanger people's safety, such as medicine, journals might have to take a more rigorous approach to verifying information as accurate.

Narayanan says that the solutions to these issues should not focus on the chatbot itself, "but rather the perverse incentives that lead to this behaviour, such as universities conducting hiring and promotion reviews by counting papers with no regard to their quality or impact".

1. Gao, C. A. *et al.* Preprint at bioRxiv https://doi.org/10.1101/2022.12.23.521610 (2022).
2. Blanco-Gonzalez, A. *et al.* Preprint at https://arxiv.org/abs/2212.08104 (2022).
3. O'Connor, S. & ChatGPT. *Nurse Educ. Pract.* **66**, 103537 (2023).