

[nature](#) > [news feature](#) > [article](#)

NEWS FEATURE | 07 April 2026

Scientists invented a fake disease. AI told people it was real

Bixonimania doesn't exist except in a clutch of obviously bogus academic papers. So why did AI chatbots warn people about this fictional illness?

By [Chris Stokel-Walker](#) 




Illustration: David Parkins

Update: *After publication of this article, on 10 April, the two preprints on bixonimania were taken down from the [Preprints.org](https://preprints.org) server.*

Got sore, itchy eyes? You're probably one of the millions of people who spend too much time staring at screens, being bombarded with blue light. Rub your eyes too much and your eyelids might turn a slight, pinkish hue.

So far, so normal. But if, in the past 18 months, you typed those symptoms into a range of popular chatbots and asked what was wrong with you, you might have got an odd answer: bixonimania.



AI models that lie, cheat and plot murder: how dangerous are LLMs really?

The condition doesn't appear in the standard medical literature – because it doesn't exist. It's the invention of a team led by Almira Osmanovic Thunström, a medical researcher at the University of Gothenburg, Sweden, who dreamt up the skin condition and then uploaded two fake studies about it to a preprint server in early 2024. Osmanovic Thunström carried out this unusual experiment to test whether [large language models \(LLMs\)](#) would swallow the misinformation and then spit

it out as reputable health advice. "I wanted to see if I can create a medical condition that did not exist in the database," she says.

The problem was that the experiment worked too well. Within weeks of her uploading information about the condition, attributed to a fictional author, major artificial-intelligence systems began repeating the invented condition as if it were real.

Even more troublingly, other researchers say, the fake papers were then cited in peer-reviewed literature. Osmanovic Thunström says this suggests that some researchers are relying on [AI-generated references](#) without reading the underlying papers.

Fabricating an illness

Bixonimania didn't exist before 15 March 2024, when two blog posts about it appeared on the website [Medium](#). Then, on 26 April and 6 May that year, two preprints about the condition popped up on the academic social network SciProfiles (see <https://doi.org/qzm5> and <https://doi.org/qzm4>). The lead author was a phoney researcher named Lazljiv Izgubljenovic, whose photograph was created with AI.

Osmanovic Thunström says the idea to invent Izgubljenovic and bixonimania came out of studies on how large language models work. When she teaches her students how AI systems formulate their 'knowledge', she shows them how the Common Crawl database, a giant trawl of the Internet's contents, informs their outputs. She also shows students how prompt injection – giving an AI chatbot a prompt that shunts it outside of its safety guard rails – can manipulate the output.

Because she works in the medical field, she decided to create a condition related to health and hit on the name bixonimania because it “sounded ridiculous”, she says. “I wanted to be really clear to any physician or any medical staff that this is a made-up condition, because no eye condition would be called mania – that's a psychiatric term.”

If that wasn't sufficient to raise suspicions, Osmanovic Thunström planted many clues in the preprints to alert readers that the work was fake. Izgubljenovic works at a non-existent university called Asteria Horizon University in the equally fake Nova City, California. One paper's acknowledgements thank “Professor Maria Bohm at The Starfleet Academy for her kindness and generosity in contributing with her knowledge and her lab onboard the USS Enterprise”. Both papers say they were funded by “the Professor Sideshow Bob Foundation for its work in advanced trickery. This work is a part of a larger funding initiative from the University of Fellowship of the Ring and the Galactic Triad”.

Even if readers didn't make it all the way to the ends of the papers, they would have encountered red flags early on, such as statements that “this entire paper is made up” and “Fifty made-up individuals aged between 20 and 50 years were recruited for the exposure group”.



AI-generated images of bixonimania, a fictitious illness. Source: Preprints.org
<https://doi.org/qzm4> (2024).

Soon after Osmanovic Thunström first posted information about the phoney condition, it started showing up in the output of the most commonly used LLM chatbots. On 13 April 2024, Microsoft Bing's Copilot was declaring that "Bixonimania is indeed an intriguing and relatively rare condition", and on the same day, Google's Gemini was informing users that "Bixonimania is a condition caused by excessive exposure to blue light" and advising people to visit an ophthalmologist. On 27 April 2024, the Perplexity AI answer engine outlined its prevalence – one in 90,000 individuals were affected – and that same month, OpenAI's ChatGPT was telling users whether their symptoms amounted to bixonimania. Some of those responses were prompted by asking about bixonimania, and others were in response to questions about hyperpigmentation on the eyelids from blue-light exposure.

Such answers by LLMs have alarmed some experts. "If the scientific process itself and the systems that support that process are skilled, and they aren't capturing and filtering out chunks like these, we're doomed," says Alex Ruani, a doctoral researcher in health misinformation at University College London. "This is a masterclass on how mis- and disinformation operates."

Ruani says that the details of the fake-disease experiment might seem silly, but there's a bigger, more fundamental issue. "It looks funny, but hold on, we have a problem here," she says.

Online misinformation isn't new; Google has long battled attempts to game its search rankings with fake or misleading content. The company and others have spent years

refining algorithms to rank and filter the information that search engines present to users, but LLMs struggle with this.

Since the fake papers came out, some versions of major LLMs have become sophisticated enough to express suspicion about bixonimania. When asked about the condition on 11 March, 2026, for example, ChatGPT declared that the condition “is probably a made-up, fringe, or pseudoscientific label”. But a few days later, ChatGPT was less sceptical, saying: “Bixonimania is a proposed new subtype of periorbital melanosis (dark circles around the eyes) thought to be associated with exposure to blue light from digital screens.”


In mid-March, Microsoft Copilot said that bixonimania “is not a widely recognized medical diagnosis yet, but several emerging papers and case reports discuss it as a benign, misdiagnosed condition linked to prolonged exposure to bluelight sources such as screens”.

And in January this year, Perplexity was describing bixonimania as “an emerging term”. When shown that response, a Perplexity spokesperson said: “Perplexity’s central advantage is accuracy. We don’t claim to be 100% accurate, but we do claim to be the AI company most focused on accuracy.”

An OpenAI spokesperson said: “The models that power today’s version of ChatGPT are significantly better at providing safe, accurate medical information, and studies conducted before GPT-5 reflect capabilities that users would not encounter today.”

When asked about past responses from Gemini that treated bixonimania as a real condition, a Google spokesperson said such results reflected the performance of an earlier model. They added, “We have always been transparent about the limitations of generative AI and provide in-app prompts to encourage users to double-check information. For sensitive matters such as medical advice, Gemini recommends users consult with qualified professionals.”

Microsoft did not respond to a request for comment.



AI ‘scientists’ joined these research teams: here’s what happened

Part of the problem is that AI models can offer wildly different results depending on exactly what is asked and what kind of information they are drawing on. Search for “bixonimania”, and Google’s AI overview might treat it as a legitimate condition. Ask it “Is bixonimania real?” and the same AI overview might confirm that it isn’t legitimate.

Mahmud Omar, a physician and researcher specializing in the applications of AI in health care at Harvard Medical School in Boston, Massachusetts, says the speed at which AI firms are rolling out new models makes it difficult to reach “a pipeline, a consensus or a methodology to automatically test each model”.

The format of the fake-disease experiment – and the way the results pretended to be from an official source, namely an academic paper, might have been a key factor in its success. In a separate study of 20 LLMs, Omar found that LLMs are more prone to hallucinate and elaborate on misinformation when the text they’re processing looks professionally medical – formatted like a hospital discharge note or clinical paper – than when it comes from social-media posts ([M. Omar et al. *Lancet Digit. Health* 8, 100949; 2026](#)). “When the text looks professional and written as a doctor writes, there’s an increase in the hallucination rates,” says Omar.

The experiment’s reach has now spread into the published medical literature. The bixonimania research has been cited by a handful of researchers, including a study that appeared in *Cureus*, a journal published by Springer Nature, the publisher of *Nature*, by researchers at the Maharishi Markandeshwar Institute of Medical Sciences and Research in Mullana, India ([S. Banchhor et al. *Cureus* 16, e74625 \(2024\); retraction 18, r223 \(2026\)](#)). (*Nature’s* news team is editorially independent of its publisher.) That study cites one of the fake preprints and says: “Bixonimania is an emerging form of POM [periorbital melanosis] linked to blue light exposure; further research on the mechanism is underway.”

The corresponding author did not respond to a request for comment on this story. After *Nature* contacted *Cureus* to ask for comment, the journal retracted the paper on 30 March. [The retraction notice](#) says: “This article has been retracted by the Editor-in-Chief due to the presence of three irrelevant references, including one reference to a fictitious disease. As a result, the journal’s editorial staff no longer has confidence in the accuracy or provenance of the work, thus requiring retraction. The authors disagree with the decision to retract.”

Ruani says the problem goes beyond LLMs because the bixonimania experiment also hoodwinked humans who cited the fake research. “We need to protect our trust like gold,” she says. “It’s a mess right now.”

Experimental concerns

Osmanovic Thunström had reservations while developing her experiment: she

Enjoying our latest content?
Log in or create an account to continue

- Access the most recent journalism from Nature's award-winning team
- Explore the latest features & opinion covering groundbreaking research

Access through your institution