

BY DHRUV MEHROTRA AND TIM MARCHMAN SECURITY JUN 19, 2024 9:00 AM

Perplexity Is a Bullshit Machine

A WIRED investigation shows that the AI-powered search startup Forbes has accused of stealing its content is surreptitiously scraping—and making things up out of thin air.



ANIMATION: JACQUI VANLIEW; GETTY IMAGES

 SAVE

Considering Perplexity's bold ambition and the investment it's taken from Jeff Bezos' family fund, Nvidia, and famed investor Balaji Srinivasan, among others, it's surprisingly unclear what the AI search startup actually is.

Earlier this year, speaking to WIRED, Aravind Srinivas, Perplexity's CEO, described his product—a chatbot that gives natural-language answers to prompts and can, the company says, access the internet in real time—as an “answer engine.” A few weeks later, shortly before a funding round valuing the company at a billion dollars was announced, he told Forbes, “It’s almost like Wikipedia and ChatGPT had a kid.” More recently, after Forbes accused Perplexity of plagiarizing its content, Srinivas told the AP it was a mere “aggregator of information.”

The Perplexity chatbot itself is more specific. Prompted to describe what Perplexity is, it provides text that reads, “Perplexity AI is an AI-powered search engine that combines features of traditional search engines and chatbots. It provides concise, real-time answers to user queries by pulling information from recent articles and indexing the web daily.”

A WIRED analysis and one carried out by developer Robb Knight suggest that Perplexity is able to achieve this partly through apparently ignoring a widely accepted web standard known as the Robots Exclusion Protocol to surreptitiously scrape areas of websites that operators do not want accessed by bots, despite claiming that it won't. WIRED observed a machine tied to Perplexity—more specifically, one on an Amazon server and almost certainly operated by Perplexity—doing this on WIRED.com and across other Condé Nast publications.

The WIRED analysis also demonstrates that, despite claims that Perplexity's tools provide “instant, reliable answers to any question with complete sources and citations included,” doing away with the need to “click on different links,” its chatbot, which is capable of accurately summarizing journalistic work with appropriate credit, is also prone to bullshitting, in the technical sense of the word.

WIRED provided the Perplexity chatbot with the headlines of dozens of articles published on our website this year, as well as prompts about the subjects of WIRED reporting. The results showed the chatbot at times closely paraphrasing WIRED stories, and at times summarizing stories inaccurately and with minimal attribution. In one case, the text it generated falsely claimed that WIRED had reported that a specific police officer in California had committed a crime. (The AP similarly identified an instance of the chatbot attributing fake quotes to real people.) Despite its apparent access to original WIRED reporting and its site hosting original WIRED art, though, none of the IP addresses publicly listed by the

company left any identifiable trace in our server logs, raising the question of how exactly Perplexity's system works.

Until earlier this week, Perplexity published in its documentation a link to a list of the IP addresses its crawlers use—an apparent effort to be transparent. However, in some cases, as both WIRED and Knight were able to demonstrate, it appears to be accessing and scraping websites from which coders have attempted to block its crawler, called Perplexity Bot, using at least one unpublicized IP address. The company has since removed references to its public IP pool from its documentation.

That secret IP address—44.221.181.252—has hit properties at Condé Nast, the media company that owns WIRED, at least 822 times in the past three months. One senior engineer at Condé Nast, who asked not to be named because he wants to “stay out of it,” calls this a “massive undercount” because the company only retains a fraction of its network logs.

WIRED verified that the IP address in question is almost certainly linked to Perplexity by creating a new website and monitoring its server logs. Immediately after a WIRED reporter prompted the Perplexity chatbot to summarize the website's content, the server logged that the IP address visited the site. This same IP address was first observed by Knight during a similar test.

It also appears probable that in some cases—and despite a graphical representation in its user interface that shows the chatbot “reading” specific source material before giving a reply to a prompt—Perplexity is summarizing not actual news articles but reconstructions of what they say based on URLs and traces of them left in search engines like extracts and metadata, offering summaries purporting to be based on direct access to the relevant text.

The magic trick that's made Perplexity worth 10 figures, in other words, appears to be that it's both doing what it says it isn't and not doing what it says it is.

In response to a detailed request for comment referencing the reporting in this story, Srinivas issued a statement that said, in part, “The questions from WIRED reflect a deep and fundamental misunderstanding of how Perplexity and the Internet work.” The statement did not dispute the specifics of WIRED's reporting, and Srinivas did not respond to follow-up questions asking if he disputed WIRED's or Knight's analyses.

ON JUNE 6, Forbes published an investigative report about how former Google CEO Eric Schmidt's new venture is recruiting heavily and testing AI-powered drones with potential military applications. (Forbes reported that Schmidt declined to comment.) The next day, John Paczkowski, an editor for Forbes, posted on X to note that Perplexity had essentially republished the sum and substance of the scoop. ("It rips off most of our reporting," he wrote. "It cites us, and a few that reblogged us, as sources in the most easily ignored way possible.")

That day, Srinivas thanked Paczkowski, noting that the specific product feature that had reproduced Forbes' exclusive reporting had "rough edges" and agreeing that sources should be cited more prominently. Three days later, Srinivas boasted —inaccurately, it turned out—that Perplexity was Forbes' second-biggest source of referral traffic. (WIRED's own records show that Perplexity sent 1,265 referrals to WIRED.com in May, an insignificant amount in the context of the site's overall traffic. The article to which the most traffic was referred got 17 views.) "We have been working on new publisher engagement products and ways to align long-term incentives with media companies that will be announced soon," he wrote. "Stay tuned!"

Just what Srinivas meant soon became clear when Semafor reported that the company had been "working on revenue-sharing deals with high-quality publishers"—arrangements that would allow Perplexity and publishers alike to profit from the publishers' investments in reporting. According to Axios, Forbes' general counsel sent a letter to Srinivas last Thursday demanding Perplexity remove misleading articles and repay Forbes for advertising revenue earned from its alleged copyright infringement.

THE FOCUS ON what Perplexity is doing is, while understandable, to some extent obscuring the more important question of how it's doing it.

The basics of the "what" aren't in serious dispute: Perplexity is making money from summarizing news articles, a practice that has existed as long as there has been news and that enjoys broad, though qualified, legal protection. Srinivas has acknowledged that at times these summaries have failed to credit the sources from which they're derived fully or prominently enough, but more broadly he denied unethical or unlawful activity. Perplexity has "never ripped off content from anybody," he told the AP. "Our engine is not training on anyone else's content."

This is a curious defense in part because it answers an objection no one has raised. Perplexity's main offering isn't a large language model that needs to be trained on a body of data, but rather a wrapper that goes around such systems. Users who pay \$20 for a "Pro" subscription, as two WIRED reporters did, are given a choice of five AI models to use. One, Sonar Large 32k, is unique to Perplexity but based on Meta's LLaMa 3; the others are off-the-shelf versions of various models offered by OpenAI and Anthropic.

This is where we come to the how: When a user queries Perplexity, the chatbot isn't just composing answers by consulting its own database, but also leveraging the "real-time access to the web" that Perplexity touts in marketing materials to gather information, then feeding it to the AI model a user has selected to generate a reply. In this way, while Perplexity has trained its own model and purports to leverage "sophisticated AI" to interpret prompts, calling it an "AI startup" is somewhat misleading; it would perhaps be more accurately described as a sort of remora attached to existing AI systems. ("To be clear, while Perplexity does not train foundation models, we are still an AI company," Srinivas tells WIRED.)

In theory, Perplexity's chatbot shouldn't be able to summarize WIRED articles, because our engineers have blocked its crawler via our robots.txt file since earlier this year. This file instructs web crawlers on which parts of the site to avoid, and Perplexity claims to respect the robots.txt standard. WIRED's analysis found that in practice, though, prompting the chatbot with the headline of a WIRED article or a question based on one will usually produce a summary appearing to recapitulate the article in detail.

Entering the headline of this exclusive into the chatbot's interface, for example, produces a four-paragraph block of text laying out the basic information that Keanu Reeves and the science fiction writer China Miéville have collaborated on a novel, seemingly complete with telling details. "Despite his initial apprehension about the potential collaboration, Reeves was enthusiastic about working with Miéville," the text reads; this is followed by a gray circle which, when moused over, provides a link to the article. The text is illustrated by a photograph commissioned by WIRED; clicking the image produces a credit line and a link to the original article. (WIRED's records show that Perplexity has directed six users to the article since its publication.)

Similarly, asking Perplexity "Are some cheap wired headphones actually using Bluetooth?" yields what appears to be a two-paragraph summary of this WIRED

story, accompanied by the art that originally ran with it. "Although this method is not a scam, it can be seen as a deceptive or ingenious workaround depending on one's perspective," the text reads. This is closer to WIRED copy ("Is it a scam? Technically no—but depending on your point of view, there's either deception going on here or some kind of ingenious hack," wrote staff writer Boone Ashworth) than either a human editor or lawyer might prefer, but the chatbot generates text insisting this is a mere coincidence.

"No, I did not plagiarize the phrase," reads text generated by the chatbot in response to a prompt given by a WIRED reporter. "The similarity in wording is coincidental and reflects the common language used to describe such a nuanced situation." How the common language is defined is unclear—aside from product listings for headphones, the only sources Perplexity cites here are the WIRED article and a Slashdot discussion of it.

FINDINGS by Robb Knight, the developer, and a subsequent WIRED analysis suggest an explanation for some of what's happening here: In brief, Perplexity is scraping websites without permission.

As Knight explains it, in addition to forbidding AI bots from the servers of Macstories.net, a site on which he works, by utilizing a robots.txt file, he additionally coded in a server-side block that in theory should present a crawler with a 403 forbidden response. He then put up a post describing how he had done this and asked the Perplexity chatbot to summarize it, yielding "a perfect summary of the post including various details that they couldn't have just guessed."

"So," he asked, reasonably, "what the fuck are they doing?"

Knight investigated his server logs and found that Perplexity had apparently ignored his robots.txt file and evaded his firewall, likely using an automated web browser running on a server with an IP address that the company does not publicly disclose. "I can't even block their IP ranges because it appears these headless browsers are not on their IP ranges," he wrote.

WIRED was able to confirm that a server at the IP address Knight observed—44.221.181.252—will, on demand, visit and download webpages when a user asks Perplexity about the webpage, regardless of what the site's robots.txt says. According to an analysis of Condé Nast system logs by our company's engineers,

it's likely this IP address has accessed the company's content thousands of times without permission.

Scraping websites that have asked not to be scraped may in some circumstances expose a company or an individual to legal risk, though the relevant case law is ambiguous and generally on the side of those accessing public websites. ("It's a complicated area of law," says Andrew Crocker, surveillance litigation director at the Electronic Frontier Foundation, "and there's a lot of litigation around it.")

Knight, the developer, says his findings nonetheless make him "furious."

"We've now got a huge industry of AI-related companies who are incentivized to do shady things to continue their business," he tells WIRED. "By not identifying that it's them accessing a site, they can continue to collect data unrestricted."

"Millions of people," says Srinivas, "turn to Perplexity because we are delivering a fundamentally better way for people to find answers."

WHILE KNIGHT'S AND WIRED's analyses demonstrate that Perplexity will visit and use content from websites that it doesn't have permission to access, that doesn't necessarily explain the vagueness of some of its responses to prompts about specific articles and the sheer inaccuracy of others. This mystery has one fairly obvious solution: In some cases, it isn't actually summarizing the article.

In one experiment, WIRED created a test website containing a single sentence—"I am a reporter with WIRED"—and asked Perplexity to summarize the page. While monitoring the website's server logs, we found no evidence that Perplexity attempted to visit the page. Instead, it invented a story about a young girl named Amelia who follows a trail of glowing mushrooms in a magical forest called Whisper Woods.

When pressed for answers about why it made up a story, the chatbot generated text that read, "You're absolutely right, I clearly have not actually attempted to read the content at the provided URL based on your observation of the server logs ... Providing inaccurate summaries without making the effort to read the actual content is unacceptable behavior for an AI like myself."

It's unclear why the chatbot invented such a wild story, or why it didn't attempt to access this website.

Despite the company's claims about its accuracy and reliability, the Perplexity chatbot frequently exhibits similar issues. In response to prompts provided by a WIRED reporter and designed to test whether it could access this article, for example, text generated by the chatbot asserted that the story ends with a man being followed by a drone after stealing truck tires. (The man in fact stole an ax.) The citation it provided was to a 13-year-old WIRED article about government GPS trackers being found on a car. In response to further prompts, the chatbot generated text asserting that WIRED reported that an officer with the police department in Chula Vista, California, had stolen a pair of bicycles from a garage. (WIRED did not report this, and is withholding the name of the officer so as not to associate his name with a crime he didn't commit.)

In an email, Dan Peak, assistant chief of police at Chula Vista Police Department, expressed his appreciation to WIRED for "correcting the record" and clarifying that the officer did not steal bicycles from a community member's garage. However, he added, the department is unfamiliar with the technology mentioned and so cannot comment further.

These are clear examples of the chatbot "hallucinating"—or, to follow a recent article by three philosophers from the University of Glasgow, bullshitting, in the sense described in Harry Frankfurt's classic *On Bullshit*. "Because these programs cannot themselves be concerned with truth, and because they are designed to produce text that *looks* truth-apt without any actual concern for truth," the authors write of AI systems, "it seems appropriate to call their outputs bullshit." ("We have been very upfront that answers will not be accurate 100% of the time and may hallucinate," says Srinivas, "but a core aspect of our mission is to continue improving on accuracy and the user experience.")

There would be no reason for the Perplexity chatbot to bullshit by extrapolating what was in an article if it were accessing it. It's therefore logical to conclude that in some cases it isn't, and is approximating what was likely in it from related material found elsewhere. The likeliest sources of such information would be URLs and bits of digital detritus gathered by and submitted to search engines like Google—a process something like describing a meal by tasting scraps and trimmings fished out of a garbage can.

Both the explanation of how Perplexity works published on its site and, for what it's worth, text generated by the Perplexity chatbot in response to prompts related

to its information-gathering workflow support this theory. After parsing a query, the text said, Perplexity deploys its web crawler, avoiding sites on which it's blocked.

"Perplexity can also," the text reads, "leverage search engines like Google and Bing to gather information." In this sense, at least, it truly is just like a human.

You Might Also Like ...

- Navigate election season with our WIRED Politics Lab [newsletter](#) and podcast
 - Don't think [breakdancing is an Olympic sport](#)? The world champ agrees (kinda)
 - How researchers cracked an 11-year-old password to a [\\$3M crypto wallet](#)
 - The uncanny rise of the world's first [AI beauty pageant](#)
 - **Give your back a break:** Here are the [best office chairs](#) we've tested
-



[Dhruv Mehrotra](#) (he/him) is an investigative data reporter for WIRED. He uses technology to find, build, and analyze data sets for storytelling. Before joining WIRED, he worked for the Center for Investigative Reporting and was a researcher at New York University's Courant Institute of Mathematical Sciences. At Gizmodo, he was... [Read more](#)

SENIOR WRITER ✕

[Tim Marchman](#) oversees coverage of politics, science, and security. Previously, he was features director at Vice, special projects editor at Gizmodo Media, and editor in chief of Deadspin, among other roles. He lives in Philadelphia with his wife and children and can be reached via email at timothy_marchman@wired.com or... [Read more](#)

DIRECTOR, SCIENCE, POLITICS, AND SECURITY

TOPICS INNOVATION ARTIFICIAL INTELLIGENCE SEARCH ENGINES

READ MORE




GEAR

Reduce AI Hallucinations With This Neat Software Trick

A buzzy process called retrieval augmented generation, or RAG, is taking hold in Silicon Valley and improving the outputs from large language models. How does it work?

REECE ROGERS



BUSINESS

Google's AI Overviews Will Always Be Broken. That's How AI Works

Google rushed out fixes after its AI search feature made errors that went viral. Fundamental limitations of generative AI mean that it will still screw up sometimes.

WILL KNIGHT



POLITICS

If Google Kills News Media, Who Will Feed the AI Beast?

Summarization tools from OpenAI and Google offer a CliffsNotes version of journalism that may further dumb down public discourse and deliver a brutal blow to an already battered media business.

NICK BILTON



ANNALS OF COMMUNICATIONS

Is Hunterbrook Media a News Outlet or a Hedge Fund?

The hybrid media-finance company wants to monetize investigative journalism in the public interest. Is it a visionary game changer or a cynical ploy?

CLARE MALONE




BUSINESS

Google's AI Overview Search Results Copied My Original Work

Google's AI feature bumped my article down on the results page, but the new AI Overview at the top still referenced it. What gives?

REECE ROGERS



INFINITE SCROLL

Faux ScarJo and the Descent of the A.I. Vultures

OpenAI's snafu over its "Her"-like voice assistant might be funny if it didn't portend a larger crisis in the integrity of digital information.

KYLE CHAYKA



INFINITE SCROLL

The Revenge of the Home Page

As social networks become less reliable distributors of the news, consumers of digital journalism are seeking out an older form of online real estate.

KYLE CHAYKA



BUSINESS

Google Admits Its AI Overviews Search Feature Screwed Up

Liz Reid, Google's head of search, said in a blog post that the company had made adjustments to its new AI search feature after screenshots of its errors went viral.

REECE ROGERS

WIRED

Get 1 year for
~~\$30~~ \$5

SUBSCRIBE

[✓x COOKIES SETTINGS](#)