



MAGAZINE • TECH

# How Aravind Srinivas turned Perplexity AI into an \$18 billion would-be Google killer

AI-powered summaries from Perplexity have pressured the search giant to adapt, and its Comet browser is an agentic AI breakthrough.



# FORTUNE



BY JEREMY KAHN  
EDITOR, AI

July 24, 2025 at 6:00 AM EDT

August/September 2025

Aravind Srinivas wanted a simple answer. If Google had given him one, things might have turned out differently—for Srinivas and for Google.

It was the fall of 2022, weeks before OpenAI would debut its viral AI chatbot, ChatGPT. Srinivas, then 28, was an ambitious AI researcher with a PhD from UC Berkeley and prestigious internships at OpenAI and Google's AI labs under his belt, plus a year working full-time for OpenAI after earning his doctorate. Now he'd left that plum job to launch a startup with three cofounders, all AI experts. Exactly what this startup would do, however, was a bit unclear.

While they were still trying to decide, Srinivas and his cofounders hired their first engineer—and that engineer needed health insurance. Srinivas had never picked an insurance provider before and had no idea how to go about it. So, like most people would, he asked Google: What is the best health insurance plan for a startup? Which HR software provides the best value for money? He quickly became frustrated with the results: The top links were all ads from insurance companies or SEO-optimized marketing content.

But a better alternative was sitting at Srinivas's fingertips. One of the business ideas he and his cofounders

had been playing around with was an AI tool that would use a large language model (LLM) to summarize information from multiple websites. The founders had even built a prototype. Srinivas remembers thinking, "Let's at least ask this tool that we have." And it started answering all these questions really well for us." Not only did the prototype help him find insurance, it also gave Srinivas his startup's mission.

Weeks later, ChatGPT launched. Seven days after that, Srinivas and his cofounders debuted their San Francisco-based company and its breakthrough product. They called both Perplexity—borrowing a machine-learning term for how well an AI model makes predictions. Even amid the buzz around ChatGPT, Perplexity gained traction among Silicon Valley's elite and, through word of mouth, among consultants, analysts, and journalists. "They all said, 'Hey, this could be something better than Google,'" Srinivas says.

While Perplexity had a search-bar-style interface that looked like ChatGPT's, its outputs were distinct and better suited to some users. ChatGPT initially had no web access, drawing answers only from its training data, with no ability to discover up-to-date information. It couldn't cite its sources and was prone to "hallucination"—in which the model confidently spewed factually wrong information. Perplexity, by contrast, searched the web, and unlike Google, which delivered blizzards of blue links to comb through, Perplexity summarized what it found. Its answers included footnotes citing web pages where the model claimed to have found those bits of information. While this method did not eliminate hallucinations, it reduced them and made them easier to spot.

ChatGPT could do almost anything involving language—including holding personal dialogues, writing code, and crafting limericks about your boss. Perplexity can do some of those things, but it's optimized to do just one thing: answer factual questions concisely and accurately. While many talked about ChatGPT being a "Google killer," the cognoscenti were whispering that the real threat was Perplexity. It could do what a Google Search could—only better. As Srinivas liked to say, he was replacing the search engine with an "answer engine."

Two and a half years later, the answer engine is worth \$18 billion—based on July reports that it raised venture capital at that valuation. Alongside prominent VC firms, the company counts Amazon founder Jeff Bezos, Google chief scientist Jeff Dean, and Meta chief AI scientist and deep learning pioneer Yann LeCun among its investors. Chipmaker Nvidia is also a backer.

Perplexity claimed \$100 million in annual recurring revenue as of March—the latest figure publicly available—earned through premium subscriptions, sponsored follow-up questions, and enterprise access to its application programming interface (API). (The startup remains unprofitable, burning cash at a rate requiring frequent venture capital injections as it fine-tunes its technology.) And it has gained a toehold—albeit a pinky toe—against Google. It now boasts close to 1 billion queries each month—impressive, though small compared with the 780 million weekly users ChatGPT claims and very small compared with Google’s 83.8 billion monthly visitors.

Google has not been killed. But the tech giant has been forced into making the most radical changes to its core search product in two decades. Google’s share of search traffic dropped below 90% at the end of 2024, according to data from Statcounter, the first time it dipped below that threshold in 15 years. While the threat from ChatGPT has been the primary driver of these traffic changes, Perplexity pioneered many of the AI search features that both Google and OpenAI have more recently rolled out.



Now Google may have to dance to Perplexity’s tune once again. The startup just launched an AI-native web browser called Comet that could represent the biggest shift in how we interact with the internet since Netscape popularized the graphical web browser in the mid-1990s.

Perplexity has become a lightning rod in debates over the relationship between AI companies and publishers, who fear that “answer engines” will obviate the need for audiences to visit their sites. But neither that controversy nor Perplexity’s financial losses have deterred tech giants from eyeing it as a juicy acquisition target as they battle for AI dominance. It has already rebuffed an approach from Meta, and it has been reported that Apple has considered bidding for the company. While some analysts doubt Perplexity can go the distance against Google and OpenAI, few question that this scrappy company is having an outsize impact on how we find information in the AI age.

**“In a world where you can easily create fake content with AI, accurate answers and trustworthy sources become even more essential.”**

ARAVIND SRINIVAS, CEO AND COFOUNDER, PERPLEXITY

Srinivas has been studying his competitor, Google, for most of his life. Growing up, he idolized current Google CEO Sundar Pichai, who hails from Srinivas’s hometown, the Indian city Chennai. During his summer internship at Google DeepMind’s London headquarters, Srinivas’s rental accommodation was so shabby that he spent many nights sleeping at the office. While exploring DeepMind’s library late one evening, he discovered *In the Plex*, a chronicle of Google’s first 15 years by journalist Steven Levy. Srinivas devoured the book, reading it multiple times. “That changed my whole understanding of startups,” he says. Reading about how two Stanford University computer-science PhD students, Larry Page and Sergey Brin, cofounded Google gave Srinivas—

halfway through his Berkeley PhD program—hope that he, too, could found a company.

At OpenAI, Srinivas was captivated by the power of LLMs to answer questions in plain English. But when he and his cofounders started Perplexity, they decided against trying to build their own LLM. “It was a decision driven through conviction and pragmatism,” Srinivas says. Pragmatism because they were broke—and training an LLM could cost tens of millions of dollars. Conviction in that it was already clear to Srinivas that AI models would become increasingly commoditized. Companies would have to differentiate themselves by what products they built on top of the base AI models, not by model capabilities alone. “There were already, like, five or six players building models,” Srinivas says. “So we thought, ‘Don’t be the yet-another-model company.’”

Srinivas and his cofounders initially used GPT-3.5, an LLM OpenAI had made available on a pay-per query basis. (A version of GPT-3.5 powered ChatGPT when it launched.) In the lingo of Silicon Valley, Srinivas’s startup would be a “wrapper”—creating a unique user interface and workflows that would “wrap around” OpenAI’s model. (Perplexity now uses a variety of models from multiple AI companies.)

As a student of Google’s history, Srinivas knew that in 2001 Page had said the ultimate version of Google would be AI providing answers, not links. But Page didn’t build that business, because AI was not yet capable enough. Instead, he and Brin built Google around links.



Srinivas now says he and his cofounders “didn’t really understand the innovator’s dilemma facing the business model for Google.” When Perplexity’s answer engine launched, Srinivas says he expected Google to copy it. “I was [thinking], ‘Okay, at some point they’re going to do it,’” he recalls. “[Sundar Pichai] keeps writing blog posts about doing it.’ But they didn’t ship it.” One reason for that: Google parent company Alphabet can’t pivot to answers without potentially jeopardizing the \$198 billion per year it makes from Google Search—a figure that represents 57% of its revenues.

Even today, three years after Perplexity’s launch, Google has not overhauled its main search landing page, google.com, to match Perplexity’s generative AI features. But it has inched ever closer: It has expanded the use of “AI Overviews,” which provide summarized capsule answers with citations above the traditional link stack. In May, it also announced a new “AI Mode” for U.S. users that functions similarly to Perplexity.

These features have started to hit Google’s link clickthrough rates, with growth slowing from high-single-digit percentages to just 2% in the first quarter of 2025, according to Bloomberg Intelligence senior technology analyst Mandeep Singh. Singh says Google’s “ad pricing remains strong, so there’s no real impact on ad revenues yet.”

Still, the landscape is shifting. “The unit of value is shifting from the click to commercial intent fulfillment,” says Chirag Dekate, a vice president, analyst at technology research firm Gartner. In other words, Google is getting paid for producing sales for e-commerce customers, not for simply driving people to a website. Dekate argues that an AI-synthesized answer that directly leads to someone making a purchase “is vastly more monetizable than a list of 10 links with ambiguous intent.”



Perplexity knows that speed is one of its advantages over Google, and it is racing to get to that next level of AI-driven commerce before the search giant. In July, the startup debuted what may be the world's first AI-native browser, Comet. Comet lets users navigate the web and use the address bar to ask Perplexity questions. But, more significantly, users can pull up an AI assistant that can read any tab the user has open and even open new ones and perform actions on the user's behalf, such as drafting and sending emails, filling out forms, booking travel, or preparing research reports—all while the user is doing other things.

Google's huge advantage in web commerce comes from knowing users' habits and preferences through Chrome, Android, Gmail, and Workspace. Perplexity sees Comet as a way to gain access to similar information. "The browser is what we live in during the day on our desktop devices. So it's just an incredibly powerful canvas," says Dmitry Shevelenko, Perplexity's chief business officer.



Dmitry Shevelenko is pursuing partnerships that could put Perplexity on more users smartphones.

WINNI WINTERMEYER FOR FORTUNE

Knowing user behavior helps Perplexity personalize its answers and makes its digital assistant more effective. The real gold mine is the assistant's ability to conduct commerce directly, potentially making Perplexity an important gatekeeper to purchases. It could act as what Microsoft's Bill Gates called a "personal agent." Whoever builds the best personal agent will dominate, Gates said, because "you'll never go to a search site again. You'll never go to a productivity tool again. You'll never go to Amazon again." Comet is Perplexity's bid to win the personal-agent race.

Perplexity isn't alone in contemplating an AI-native web browser. OpenAI's browser is expected to launch imminently, and in July, the company rolled out an AI agent that works independently of a browser. Google, meanwhile, could update Chrome to make it more agentic. But Shevelenko argues that with 3 billion people using Chrome, implementing an agentic AI browser—which uses vastly more computing resources than a conventional browser—could be too expensive for Google to roll out broadly. Google's own mass, in essence, is now constraining it.

"Trust." The term comes up a lot in talking to Perplexity executives. "There is going to be a fundamental re-architecture of the internet around giving people answers instead of links," Shevelenko says. "How do you trust that answer? There is going to be a generational company born out of that."

Srinivas likes to talk about trust, too, seeing it as Perplexity's fundamental differentiator.

“In a world where you can easily create fake content with AI, accurate answers, trustworthy sources become even more essential,” he says. He says that other AI companies would struggle to match Perplexity’s factuality because those companies had competing goals. ChatGPT needs to be able to write poetry and brainstorm marketing ideas, give you advice about how to come out to your parents, and tell you how a nuclear reactor works. The creativity required for some of these tasks means OpenAI can never fully optimize ChatGPT for factual accuracy. For OpenAI, “hallucination is a feature,” Srinivas says. For Perplexity, “hallucination is a bug.”

Because Perplexity needs to surface accurate, up-to-date information, its relationship with news publishers is vital—but also fraught. Ask Perplexity a question about current events and its web crawlers will search news sites, scraping their content to inform its answer. Many publishers see this as an existential threat, disintermediating their relationship with their audience and denying them the ability to monetize traffic through advertising and subscriptions.

Jessica Chan, who heads Perplexity’s publishing partnerships, says she has fought to overcome these fears. “Our success is tied to a thriving journalism and digital publishing ecosystem, because we know these journalists produce these high-quality, verified facts,” Chan says. “We need the continual production of that type of information. There is really no world in which Perplexity is successful, but publishers are not.”



Jessica Chan is working to convince publishers that Perplexity is an ally.

WINNI WINTERMEYER FOR FORTUNE

Perplexity has partnered with several publishers, including *Time*, *Le Monde*, *Der Spiegel*, and the *Los Angeles Times* (and, full disclosure, *Fortune*). The company shares money it earns from sponsored follow-up questions that appear beneath AI-generated answers sourced from a partner publication’s content. Perplexity also gives these publications access to its enterprise product and helps them build their own AI applications.

Still, Perplexity's "answer engine" has angered plenty of publishers. Some accuse the company of flagrantly violating "robots.txt," a voluntary protocol publishers use to signal that their content is off-limits to bots, despite Perplexity saying it would abide by the standard. Forbes threatened legal action against Perplexity for both scraping its content without permission and allegedly plagiarizing it verbatim in AI-generated answers without adequate citation. The BBC has sent a similar cease-and-desist letter to the company. And News Corp.—owned DowJones and the New York Post have sued Perplexity for copyright violations as well as "trademark dilution" for attributing inaccurate AI-generated answers to News Corp. Perplexity has said it was "surprised and disappointed" by News Corp.'s suit and called its allegations "misleading at best."

Srinivas says clashes are unfortunate, especially since Perplexity has been more upfront than Google about AI's future: "We've been very transparent that this new interface will not send you as much traffic anymore."

Data from internet infrastructure provider Cloudflare support this: In 2015, Google scraped publisher websites twice for every person who clicked a link from Google to visit that site (which is called a referral). Recently, that's increased to 18 times per referral. For AI chatbots, the ratio is hugely lopsided: OpenAI scrapes 1,500 times per referral, Anthropic 60,000 times. Cloudflare, which handles 20% of internet traffic, has now started blocking web crawlers like the ones Perplexity uses by default, giving publishers the option to "white-list" some or to get paid per crawl. Chan says Perplexity applauds Cloudflare for experimenting with ways to help publishers find a business model that works for the AI era.

Perplexity pioneered many features that have since been copied by rivals, from citations to follow up questions. It was the first company to deploy DeepSeek's R1 "reasoning model" and show users the model's chain of thought—the AI's internal dialogue about what steps it plans to take and why. "There's a lot of things that we've done, micro-innovations around how we present answers, and our follow-up questions—a lot of stuff that has shown up in other products, just because it works," says Henry Modisett, Perplexity's head of design.



Henry Modisett, VP of Design at Perplexity AI.  
WINNI WINTERMEYER FOR FORTUNE

Perplexity's biggest challenge is brand awareness. Everyone knows Google. Most people have played around with ChatGPT. Only a select few have heard of Perplexity. It's why some analysts are skeptical Perplexity's Comet will be a breakout hit. "Consumers still trust Google's information more than ChatGPT's or Perplexity's," says Nikhil Lai, an analyst at tech research firm Forrester.

Shevelenko acknowledges that the company has spent little on marketing so far, but he says that it plans to ramp up advertising around Comet. Perplexity is also targeting smartphone makers to gain wider distribution. It inked a deal with Motorola to have Perplexity preinstalled on its Razr and Edge 60 devices, and new customers will get a free three-month subscription to Perplexity Pro (which normally costs \$20 per month). It announced a similar tie-up with Indian telecom carrier Airtel. Perplexity has also been in talks with Samsung about a partnership that would see Perplexity's app preloaded onto Samsung's top-of-the-line Galaxy smartphones and perhaps integrated into Samsung's mobile browser. It could also see Perplexity displace Google Gemini as the AI brain behind Samsung's digital assistant, Bixby.

The U.S. government may also provide a critical tailwind. Last year, a federal district court judge ruled that Google maintains an illegal monopoly in search, and he is currently contemplating what remedies to order. The U.S. Justice Department has asked the judge to force Google to spin off Chrome and ban Google from paying to position its search engine as the default in others' products. Srinivas, who was the only AI startup CEO to testify in the remedies portion of the trial, says he thinks Google should retain Chrome but be compelled to allow Android users to more easily select alternative browsers—including, say, Comet.

Whatever the judge decides, Shevelenko says that the antitrust verdict has already helped the startup. "A lot of [phone makers] and carriers are more willing to work with us, just because they feel Google is going to be less aggressive in penalizing them or punishing them for engaging with a competitor," he says.

Srinivas has heard these kinds of doubts before. In some ways, he's been proving doubters wrong his whole life. He says he grew up in a family that was "lower-middle-class, even by Indian standards." His father worked as a chartered accountant; his mother worked for India's equivalent of the Social Security Administration. He studied electrical engineering, not computer science, as an undergrad. "I had no network," he says. "I had no mentors."

Yet here he is, now commanding some of tech's most prominent stages. On a sweltering June afternoon, Srinivas, wearing a pale linen suit and lavender shirt, is speaking in a grand 19th-century hall in Oxford, England. The hall was purpose-built to host debates of the Oxford Union, the university's august debating society, and appearing here has become *de rigueur* for an American tech CEO on the make. It's a way to get on the radar of potential recruits—especially computer-science students—and to generate headlines and social media buzz.

Srinivas tells the assembled Oxonians that despite being surrounded by hundreds of years of history, they should not become hidebound by tradition. "AI will not be consumed with tradition," he says. "There have been many times in history when the conventional wisdom was wrong."

Srinivas is hoping this is one of those times.

*This article appears in the August/September 2025 issue of Fortune with the headline "Is this man a Google killer?"*