

In-depth insights in seconds. Ask Deep Research. →

AI Agenda

OpenAI's Math Breakthroughs Rattle Cryptography Leaders



Art via Getty Images



By Tiffany Li

Oct 14, 2026, 7:00am PDT



OpenAI's release last Tuesday of more than 300 math results it says an internal model solved has sparked concern among some of the world's leading cryptographers, or experts in encoding information to keep it private and secure. They worry that AI's continued math advances could find new algorithms that make our current methods of protecting information on encrypted services, from online payments to messaging platforms like WhatsApp, easier to crack.

The basis for modern internet security is the assumption that some math operations are easy to do, such as a process of encoding using a number known as the public key. But they are hard to undo, or decode, without knowing a specific number, called the private key. This assumption of difficulty is the foundation of secure communication over the internet, which has helped enable e-commerce, cryptocurrencies and digital accounts.

AI models from OpenAI and others, which have been solving problems that have stumped math experts for decades, have raised the prospect that powerful AI systems could find ways to solve those problems without that key.

You're reading the **AI Agenda** Newsletter

[View all newsletters](#)

“I am anticipating a vast change in mathematics, cryptanalysis included,” said **Whitfield Diffie**, who helped invent public key cryptography in 1976 with the **Diffie-Hellman key exchange algorithm**, in an email to me.

“I don’t know the timescale. A year is probably too soon; five may not be,” said Diffie, now a consulting scholar at the **Stanford Center for International Security and Cooperation**.

Martin Hellman, a professor emeritus of electrical engineering at **Stanford University** who created public key encryption with Diffie, was more optimistic. He said he expected AI-aided attacks to disrupt cybersecurity in the short term. In the long term though, he expected people to better integrate AI as a tool to patch security holes.

Such comments were more mild than some that rang out on X last week. “I think we might lose public key cryptography,” **Matthew Green**, a professor of computer science at **Johns Hopkins University** and founder of encryption startup **Zeutro**, wrote after OpenAI’s math release.

And **Justin Drake**, a security researcher at the **Ethereum Foundation**, called upon the blockchain industry to plan for “bunker mode” due to possible AI advances in cryptocurrency.

OpenAI's Crypto Efforts

Possibly increasing the alarm was the fact that OpenAI's trove of math papers did not include any cryptography results. The absence caused some experts to speculate OpenAI was sitting on results it wasn't sharing publicly.

OpenAI may have chosen to disclose any cryptography discoveries to the U.S. government first because of their security importance, Green told The Information in an interview. There's a chance the public might never find out about any such discoveries, says **Hed Kovetz**, CEO of cybersecurity firm **Silverfort**.

OpenAI did not comment on whether it has tried to use internal models to solve cryptographic problems.

Green said he believes AI companies have already put substantial resources into solving cryptography problems. Indeed, OpenAI has previously released theoretical results in quantum complexity and lattice cryptography it said were found by its latest model **Astra**, released in September. OpenAI has also employed prominent cryptography experts including **Harvard University** professor **Boaz Barak**, who has made breakthroughs on the limits of secrecy in cryptography and is a member of OpenAI's technical staff.

Green said he could imagine frontier AI in the next two years making a more fundamental break in elliptic curve cryptography, one of a small number of common encryption algorithms.

"Things get really bad if that happens because it doesn't have to happen across too many problems, it just has to happen across two [algorithms]" to impact a large share of encryption, Green said. A single elliptic curve encryption algorithm called X25519 "now secures the majority of Internet connections," web infrastructure and security company **Cloudflare** found in 2025.

The loss of elliptic curve cryptography could necessitate a shift to less understood encryption methods that might be even less secure, Green said.

Others are more skeptical of the possibility of AI breaking common encryption methods.

"These are problems that we designed specifically for them being not breakable," said **Ewin Tang**, a professor of computer science at **Princeton University**. Tang became something of a legend in math at age 18, when she disproved decades of quantum theory by finding classical algorithms that match the speed of quantum algorithms.

Tang believes the implementation of the algorithms into code is likely to contain more vulnerabilities than the algorithms themselves. Still, while she expects widely-studied common algorithms to largely be safe, Tang said some less-studied encryption algorithms, including those meant to be unsolvable by quantum computers, could be solved or weakened by AI.

Anthropic had already released results this July that largely broke one post-quantum algorithm, which had been proposed to the **National Institute of Standards and Technology**. Tang said more of the same could follow.

Anthropic shared its July research in advance with the U.S. government and industry partners and plans to continue to disclose cryptography work, a spokesperson for Anthropic told The Information. “We will responsibly disclose and act on any significant findings if and when we find anything,” the spokesperson wrote.

Here's what else is going on...

OpenAI's Crime Thriller Moment

The news Tuesday that the FBI had arrested a Saudi national working at OpenAI had all the makings of a crime thriller—or at least the pilot of one! The employee, a Saudi national who had previously worked at **Meta**, had allegedly inserted what **OpenAI** called a “sophisticated recording device,” wrapped in toilet paper, into the USB port in the back of a TV in a conference room used by its executives, the Justice Department said.

But while the news spooked OpenAI employees, some very senior ones had already been girding for such attacks.

Well before **Dario Amodei** left OpenAI to start **Anthropic** in 2021, he had pushed OpenAI to review research for potential “information hazard” risks that could potentially tip off U.S. adversaries like China, my colleagues reported in September. Amodei worried that China, Russia or other nations could hack into OpenAI and steal its research. In fact, he often wrote memos at home on an air-gapped computer, meaning one isolated from Wi-Fi and external networks, and printed out the memos to share with colleagues, rather than work in Google Docs. Such precautions don't look so odd anymore!

Overheard

Microsoft has discussed making significant cost reductions in its **Azure** cloud unit, *The Information* reported late Tuesday. **Girish Bablani**, president of Microsoft's Azure cloud unit, recently told senior leaders in his group to think about ideas for reducing costs by at least 20%.

JPMorgan CFO **Jeremy Barnum** said on Tuesday that the bank is carefully monitoring its AI exposure as the capital expenditure demands of the AI buildout grow. JPMorgan also said it's blocking external AI agents like **Meta's Muse** and **Instinct** from accessing client accounts.

Character.ai, a five-year-old startup that lets users create and interact with AI-powered personas, on Monday said it filed for Chapter 11 bankruptcy after encountering "significant legal, regulatory, and operational headwinds."

Court Watch

A judge has ruled that **Apple** will gain expedited access to **OpenAI's** hiring materials in the lawsuit Apple filed against the ChatGPT maker in July, alleging that it was stealing trade secrets. Apple will be able to view documents, like interview checklists, guides, instructions, manuals and scripts, associated with OpenAI's interviews with former Apple employees.

Deals and Debuts

See *The Information's* [Generative AI Database](#) for an exclusive list of private companies and their investors.

DayOne, a Singapore-based data center developer, raised about \$3.67 billion, including a roughly \$2.4 billion offshore term loan facility, from seven banks to fund its expansion into Malaysia for AI workloads.

Infinigence AI, which builds AI computing infrastructure, raised 2.3 billion yuan (about \$344 million) in a Series C funding round from investors including **Lenovo**, **Tencent**, **Baidu**, and several state-backed funds.

BlueCore Compute, which designs AI server chips, raised nearly 2 billion yuan (about \$280 million) in a Series B funding round from investors including **Shenzhen Capital Group** and several government-backed funds.

ICON, which 3D-prints buildings with robots, raised \$200 million in a Series D funding round led by **Baillie Gifford** and funds managed by **BlackRock**.

Applied Intuition, which makes simulation software for self-driving vehicles, has received investment offers that would value it at more than \$30 billion, up from \$15 billion in June, The Information reported.

Atribot, also known as **Stardust Intelligence**, which builds humanoid robots, raised 1 billion yuan (about \$149 million) at a 15 billion yuan (about \$2.2 billion) valuation in a funding round led by **Shenzhen Investment Holdings**-linked capital.

Mosi Intelligence, which builds multimodal AI models, raised more than 1 billion yuan (about \$141 million) in a Pre-Series A funding round led by **Huakong Fund**, with participation from **China Development Bank**-linked capital.

Suki, which uses AI to help clinicians document patient visits, raised \$90 million in Series E funding from **Healthier Capital**, **Crosslink Capital**, **Hercules Capital** and existing investors.

Albert Invent, which makes an AI platform for chemistry research, raised \$55 million in a Series B funding round led by **TCV**.

Reach Security, which finds gaps in corporate cyber defenses, raised more than \$50 million in a Series B funding round led by **Top Tier Capital Partners**.

Pluto Security, which monitors AI agents for security risks, emerged from stealth with \$42 million in total funding, including a \$34 million Series A funding round led by **Jump Capital**.

Scuba, which gives teams and AI agents shared memory, raised \$34 million across two funding rounds: a pre-seed led by **Accel** and a seed led by **Accel** and **IVP**.

Duvo, a Prague, Czech Republic and New York City-based startup that develops AI to automate business operations, raised \$32 million in Series A funding led by **Balderton Capital**.

Scimplify, which uses AI to produce specialty chemicals and critical materials, raised \$25 million in funding led by **Hitachi Ventures**.

Brellium, which audits medical billing for compliance using AI, raised \$23 million in a funding round led by **Left Lane Capital** and **First Round Capital**.

Causa Prima, which uses AI to settle payment disputes, raised \$20 million in a seed funding round led by **Base10 Partners**.

Undecimal Bio, which uses AI to design RNA-drug delivery molecules, raised \$18 million in a seed funding round led by **4BIO Capital** and **IQ Capital**.

Mogul, which uses AI for real-estate investing, raised \$15.5 million in a Series A funding round led by **Draper Associates** at a \$126 million valuation.

Ubitium, which makes chips for robots and drones, raised \$14 million in a seed funding round led by **Ibex Investors**.

Ghost, which makes a computer for local AI agents, raised \$11 million in a seed funding round led by **Andreessen Horowitz**.

Clueso, which uses AI to make training videos, raised \$10 million in a Series A funding round led by **Dharana Capital**.

Harbor, which insures companies against AI agent risks, raised \$10 million in a seed funding round led by **Base10 Partners**, **Abstract** and **8VC**.

Onyx Bio, which uses AI to model immune responses, raised \$8 million in a seed funding round led by **Dimension Capital**.

CarbonFarm, which uses AI to monitor farm methane emissions, raised €7 million (about \$7.6 million) in a Series A funding round led by **Colam Impact**.

Collective X Health, which builds an AI platform for women's health, emerged from stealth with \$3.2 million in pre-seed funding led by **Council Capital** and **Anthemis' Female Innovators Lab Fund**.

Decagon, which makes AI customer-service agents and was last valued at \$4.5 billion, has hired a bank and is in early talks with potential buyers, including **ServiceNow**, Business Insider reported.

OpenAI is in early talks to acquire **Trail of Bits**, which finds and fixes vulnerabilities in open-source software, The Information reported.


Anthropic has struck an agreement to license technology and hire some engineers from **Opaque Systems**, a five-year-old startup that sells software for securing confidential data for use with AI agents, The Information reported.

Tencent raised about \$5.7 billion through US dollar and offshore yuan bond offerings, drawing more than \$23 billion in orders, mainly to fund AI capital expenditure such as chip purchases.

Jabil and **Adani Enterprises** signed a non-binding agreement for a proposed 50-50 joint venture to manufacture AI servers, racks and data center equipment in India, aimed at both domestic demand and exports.

Apple is launching a tablet-like smart home device as part of a renewed push into the home. The \$299 **HomeView** is a 7-inch screen that sits on kitchen counters or mounts on walls.

New From Our Reporters



Exclusive


OpenAI Scrambles to Fix Safety Issues Before Shipping Its Next Flagship Model

By Stephanie Palazzolo, Aaron Holmes and Amir Efrati



Inside a TPG-Backed Accounting Firm's Plans for AI and More Acquisitions

By Alex Eichenstein



Exclusive

OpenAI Eyes Acquisition That Could Boost Cybersecurity Business

By Valida Pau and Amir Efrati



Exclusive

Microsoft Cloud President Asks Leaders to Find 20% Cost Cuts

By Kevin McLaughlin and Aaron Holmes

Recommended Newsletter

Every weekday, The Briefing helps executives get smarter about the latest in tech, media and finance. [Subscribe for free now.](#)

What We're Reading

AI Chips' New Secret Sauce Is Packaging

An Open Letter on AI and Cybersecurity from the Scientific Community

Open models have real benefits. Do they outweigh the risks?

Upcoming Events

Tuesday, October 27 — The Information's 2026 WTF Summit

The Information's WTF Summit returns to Napa Valley, bringing together the women shaping tech, media and finance. Confirmed speakers include Whitney Wolfe Herd (Bumble), Anne Wojcicki (23andMe Research Institute), Kim Posnett (Goldman Sachs) and comedian & entrepreneur Whitney Cummings, with more to come. Expect candid conversations and a room built for meaningful exchange. Tickets are selling fast—reserve your spot now.

[More details](#)

Tuesday, November 10 — The Long Road to an IPO

The Information's Dakin Campbell will host a panel discussion on the evolving landscape of public listings and exit opportunities during SF Tech Week, sponsored by NYSE.

[More details](#)

Thank you for reading the AI Agenda Newsletter! I'd love your feedback, ideas and tips: stephanie@theinformation.com.

If you think someone else might enjoy this newsletter, please pass it forward or they can sign up [here](#).

Tiffany Li is a reporter at The Information's San Francisco office. She is the recipient of a Tarbell Fellowship, funded by the Tarbell Center for AI Journalism, which provides training in AI and journalism. She previously covered politics and

campaigns for ABC News and WisPolitics. She also served as the editor-in-chief of the Chicago Maroon, the University of Chicago's student newspaper. You can reach her email at tiffany.li@theinformation.com or Signal at [@tiffanyli.99](https://tiffanyli.99).