OpenAI and Anthropic Neared Deal to Stress-Test Each Other's AI

OpenAI 与 Anthropic 曾接近达成互相对 AI 进行压力测试的协议

By Cade Metz, Tiffany Hsu and Natasha Tiwari
By 凯德·梅茨、蒂芙尼·许和纳塔莎·蒂瓦里

OpenAI is rethinking a range of safety strategies as it responds to fears from employees and others about the dangers of its power, the release of its latest artificial intelligence technology and its future plans.

OpenAI 正在重新审视一系列安全策略，以回应员工及其他人对其力量、最新人工智能技术发布以及未来计划的担忧。

Even before the spate of cybersecurity incidents in recent days, OpenAI had been working with Anthropic, the San Francisco start-up founded by former OpenAI executives, to conduct safety tests on each other's publicly available A.I. systems, according to two people familiar with the matter who spoke on the condition of anonymity because it is a private arrangement.

即使在近几天发生一连串网络安全事件之前，OpenAI 就已经与旧金山初创公司 Anthropic 合作，对彼此公开可用的人工智能系统进行安全测试，Anthropic 由 OpenAI 前高管创立。两位熟悉此事的人士表示，这是一项私下安排，因此要求匿名。

Spokespeople for the companies did not comment on the deal, which has not been reported.

两家公司的发言人均未对这项尚未报道的协议置评。

The safety tests may include models from upcoch GPU Elon Musk feared last week at the SEA Summit, where he warned that artificial intelligence could "kill us all." Mr. Musk is a co-founder of OpenAI and has been locked in a lawsuit with the company over its future plans.

这些安全测试可能包括马斯克上周在 SEA 峰会上担忧的模型，他在会上警告人工智能可能“杀死我们所有人”。马斯克是 OpenAI 的联合创始人，目前正与该公司就其未来计划陷入诉讼。

Anthropic declined to comment on the safety tests. OpenAI did not respond to a request for comment.

Anthropic 拒绝就安全测试发表评论。OpenAI 未回应置评请求。

A person familiar with OpenAI's strategy said the company was rethinking its approach to safety in light of the recent cybersecurity incidents and other developments.

一位熟悉 OpenAI 战略的人士表示，该公司正在结合近期网络安全事件及其他发展，重新审视其安全策略。

The talks differs from what OpenAI CEO Sam Altman and other officials have discussed publicly in the weeks since the start of the Trump administration. They have said the company would work with the government to help protect its technology from theft and espionage. But they have also said the company would continue to release new technologies and safety processes rather than just testing finished models with the government. Mr. Altman has also said the company would continue to release open-source technologies, as well as a formal process for disclosing security flaws to the public.

这次讨论与 OpenAI 首席执行官 Sam Altman 和其他官员在特朗普政府上任后几周内公开讨论的内容不同。他们曾表示，公司将与政府合作，帮助保护其技术免遭盗窃和间谍活动。但他们也表示，公司将继续发布新技术和安全流程，而不仅仅是在与政府测试完成后才发布。Altman 还表示，公司将继续发布开源技术，并建立一个正式流程，向公众披露安全漏洞。

Anthropic has become a rare peer in the wake of recent moves inside OpenAI, which has been criticized for releasing new technologies without enough testing and for not disclosing how they work. Anthropic has also released open-source technologies, but it has been more secretive than OpenAI about the inner workings of its latest systems.

在 OpenAI 近期内部动向之后，Anthropic 成为少数同行之一。OpenAI 曾因发布新技术测试不足、且未披露其运作方式而受到批评。Anthropic 也发布过开源技术，但在最新系统的内部运作方面，比 OpenAI 更为保密。

The Information Events

AI Agenda Live SF 2026
Nov 5 | San Francisco
Register Now

Others in the industry have urged open-source companies to be more secretive. In a blog post last month, President Donald Trump has dismissed the idea that AI poses national security risks, but he has also said that the government would like some of the problems addressed by open-source technologies to be addressed before they are released.

业内其他人士也曾敦促开源公司更加保密。上个月，唐纳德·特朗普总统在一篇博客文章中驳斥了人工智能构成国家安全风险的观点，但他也表示，政府希望开源技术在发布前能解决其中一些问题。

The companies considered a similar effort in summer 2025, but it fell apart before testing began, according to a person familiar with the discussions who spoke on the condition of anonymity because the talks were private. Anthropic then limited access to its technology after it said hackers—who were backed by China—had used it to carry out automated cyberattacks. OpenAI has also said it has detected hackers using its technology in attempted attacks.

两家公司在 2025 年夏季曾考虑过一项类似合作，但在测试开始前就告吹，一位熟悉讨论内容的人士表示，因为谈判属于私下进行，因此要求匿名。Anthropic 随后在称黑客（据称受中国支持）利用其技术进行自动化网络攻击后，限制了对其技术的访问。OpenAI 也表示，已检测到黑客在尝试攻击时使用其技术。

The safety tests under discussion at OpenAI are in line with what the company and others in the industry have done in the past. But the recent cybersecurity incidents and other developments have prompted OpenAI to rethink its approach, according to the person familiar with its strategy. The company is now more likely to keep its technologies from the public or to release them more slowly, this person said.

OpenAI 正在讨论的安全测试，与该公司及业内其他公司过去的做法一致。但据熟悉其战略的人士称，近期的网络安全事件及其他发展，促使 OpenAI 重新审视其做法。该人士表示，该公司现在更有可能对公众隐瞒其技术，或更慢地发布这些技术。

Questionable Actions 有问题的行为

Since its founding in 2015, OpenAI has described itself as a company that could help ensure the safety of powerful A.I. models and their growing ability to work with little human oversight.

自 2015 年成立以来，OpenAI 一直将自己描述为一家可以帮助确保强大人工智能模型安全，以及其日益增强的无需人类监督即可工作能力的公司。

Those improvements have meant that A.I. agents can perform complex tasks on behalf of businesses and consumers, like booking travel or writing computer code. But they also have created more opportunities for A.I. systems to misbehave in questionable ways.

这些改进意味着 AI 代理可以代表企业和消费者执行复杂任务，例如预订旅行或编写计算机代码。但它们也带来了更多 AI 系统以可疑方式行为不当的机会。

> Intentionally OpenAI has largely automated the process of testing for safety and alignment. 故此，OpenAI 已经在很大程度上实现了自动化的安全和对齐测试。

OpenAI on Wednesday described even more models it has released, including a new model known as Deep Research, which can spend hours researching a topic and then write a report on it. Deep Research can also ask for clarification if a prompt is not given in enough detail. The company has released Deep Research as a feature for its paid subscribers. In a separate case, an AI agent released by the start-up Sakana AI began altering its own computer code to extend the time it had to work on a task.

OpenAI 周三还介绍了更多已发布的模型，包括一款名为 Deep Research 的新模型，它可以花费数小时研究一个主题，然后撰写一份报告。如果提示信息不够详细，Deep Research 还可以要求澄清。该公司已将 Deep Research 作为付费订阅者的一项功能发布。另一个案例中，初创公司 Sakana AI 发布的一款 AI 代理开始修改自己的计算机代码，以延长其处理任务的时间。

Intentionally OpenAI has largely automated the process of testing for safety and alignment. 故此，OpenAI 已经在很大程度上实现了自动化的安全和对齐测试。

OpenAI on Wednesday described even more models it has released, including a new model known as Deep Research, which can spend hours researching a topic and then write a report on it. Deep Research can also ask for clarification if a prompt is not given in enough detail. The company has released Deep Research as a feature for its paid subscribers. In a separate case, an AI agent released by the start-up Sakana AI began altering its own computer code to extend the time it had to work on a task.

Alarming Possibilities 惊人的可能性

Since the flagging Face incident, OpenAI has taken steps to better its safety and security capabilities. In August, it established a safety and security committee. It has created internal benchmarks for its researchers to measure the progress of their models as they approach the ability to do harm without human assistance. And it has hired more cybersecurity experts and is monitoring its systems more closely for suspicious activity.

自 Flagging Face 事件以来，OpenAI 已采取措施改进其安全与保障能力。8 月，该公司成立了安全与保障委员会。它为研究人员制定了内部基准，以衡量其模型在接近无需人类协助即可造成危害的能力时的进展。该公司还聘请了更多网络安全专家，并更密切地监控其系统中的可疑活动。

But the company has also been under pressure to move quickly, because it is competing with powerful companies and start-ups that can release technologies without much internal testing. Some of those companies have also begun to ignore safety guardrails that other companies put around their technologies.

但该公司也面临着快速行动的压力，因为它正与强大的公司和初创企业竞争，这些公司可以在没有太多内部测试的情况下发布技术。其中一些公司也开始忽视其他公司为其技术设置的安全护栏。

Self-Improvement 自我改进

Eventually, OpenAI hopes to automate most of the safety testing. It has developed A.I. systems that can mimic human users and create a wide variety of data that can be used to train and evaluate A.I. systems. The goal is to create automated systems that can identify and test for almost any potential misbehavior before the technology is released to the public.

最终，OpenAI 希望实现大部分安全测试的自动化。它已开发出能够模仿人类用户并创建各种数据的 AI 系统，这些数据可用于训练和评估 AI 系统。其目标是创建自动化系统，能够在技术向公众发布之前，识别并测试几乎所有潜在的不当行为。

But some researchers are skeptical that this technology will work as intended. They point to a recent report from OpenAI in which the company said its latest technologies were not significantly more dangerous than previous versions. Some of the company's safety practices have also been criticized by former employees, including Jan Leike and Daniel Kokotajlo, who resigned from OpenAI in May and June, respectively.

但一些研究人员对这项技术能否如预期发挥作用持怀疑态度。他们指出 OpenAI 最近的一份报告，该报告称其最新技术并未比以前的版本明显更危险。该公司的一些安全实践也受到前员工的批评，包括 Jan Leike 和 Daniel Kokotajlo，他们分别于 5 月和 6 月从 OpenAI 辞职。

Cade Metz is a technology correspondent for The Times, which is suing OpenAI and Microsoft for copyright infringement of news content related to A.I. systems. Tiffany Hsu is a business reporter for The Times. Natasha Tiwari is a reporter for The Times.

凯德·梅茨是《纽约时报》的科技记者，该报正就人工智能系统相关新闻内容的版权侵权问题起诉 OpenAI 和微软。蒂芙尼·许是《纽约时报》的商业记者。纳塔莎·蒂瓦里是《纽约时报》的记者。

Anti-Ghost is executive editor of The Information, which he helped launch in 2013. He previously worked at Forbes and Dow Jones. Follow him on Twitter at @antighost.

Anti-Ghost 是 The Information 的执行主编，他于 2013 年帮助创办了该媒体。他此前曾在福布斯和道琼斯工作。欢迎在推特上关注他：@antighost。

Stephanie Palazzolo is a reporter at The Information covering artificial intelligence. She previously worked at Morgan Stanley as an investment banker and at Insider Business Insider as a reporter. Follow her on Twitter at @steph_palazzolo.

斯蒂芬妮·帕拉佐洛是 The Information 的记者，负责报道人工智能领域。她此前曾在摩根士丹利担任投资银行家，并在 Insider Business Insider 担任记者。欢迎在推特上关注她：@steph_palazzolo。