



Written evidence submission to the Business, Innovation, Science and Trade Select Committee
5 October 2026

OpenAI is a leading artificial intelligence research and deployment company. Our mission is to ensure that artificial general intelligence benefits everyone. In service of our mission, we research, develop, and release cutting-edge AI technology as well as tools and best practices for the safety, alignment, and governance of AI. Our goal is to build AI systems that help solve hard problems, in order to contribute to a world where everyone benefits from the social, economic, and technological opportunities of AI.

We are grateful for the opportunity to appear at the Committee's oral evidence session on 13 October. This session comes at an important time: the UK has an opportunity to play a leading role in the next chapter of AI in innovation and safety. We have called for both international standards and stronger rules on frontier AI safety. We welcome the Committee's consideration of how the UK can turn that ambition into action.

We are submitting written evidence in advance of this session, focusing on how AI capabilities are advancing, the opportunity this presents for the UK, and how we approach safety and regulation.

1. AI is entering a new chapter.

AI capabilities are advancing rapidly, including the ability to support increasingly complex scientific and technical work. These developments could accelerate discoveries in medicine, support scientific discovery, and expand economic opportunity.

For example, our models recently solved the [Navier–Stokes problem](#), widely considered one of the deepest problems at the frontier of mathematics, whose equations are used to model how fluid moves and have direct applications for aircraft design, weather forecasting, and the study of blood flow. We pursue advances like this not only for mathematics sake, but for their potential to help researchers better understand and ultimately improve human health, and to bring other significant benefits to society.

In this new era, we're working to ensure that AI works for everyone, everywhere. The measure of success should be how widely AI's benefits are shared, giving people across every postcode in Britain the tools and skills to learn, create, work and solve problems.

2. Britain is exceptionally well placed to seize this opportunity.

A new chapter in AI capabilities demands a new chapter in AI policy. Britain's strengths in science, talent and enterprise can translate AI into new discoveries, growing businesses, better services and good jobs.

Millions of people across the UK use ChatGPT daily for work, education and everyday life. The UK is one of the top 10 countries globally for weekly active ChatGPT users and one of our top five for paying subscribers and API developers. Every day, people and organisations, from developers and startups to leading enterprises, are using our technology to build, innovate and grow.

The opportunity now is to turn that progress into practical benefits for more people: helping them in their daily lives, improving the public services they depend on, and accelerating scientific discovery.



That is why OpenAI is building here for the long term. London was our first international office and will become our largest international research hub. Through this investment, we will help broaden access, skills and adoption across the country, supporting research and talent, accessible AI training, and adoption across businesses, education and public services so benefits reach across the UK.

3. Our commitment to safety.

As capabilities grow, safety must increasingly set the pace of AI progress. Safety does not stand in the way of progress; it is what allows progress to go further and benefit more people.

We have strengthened monitoring, alignment, and security safeguards across the model-development lifecycle, including stronger isolation for frontier research workloads, expanded monitoring of model behaviour during tool-enabled training and evaluations, and clearer rules for when to escalate concerns. For our latest model Astra, we also introduced monitoring of full trajectories, including chains of thought, and a mandatory alignment-evaluation gate before broader internal deployment.

As capabilities advance, our safeguards must keep pace. In practice, this can mean pausing training to strengthen safeguards and make additional alignment improvements before training resumes. For example, we have previously paused training of certain frontier models while we strengthened our training infrastructure, monitoring, and approach to alignment. Following recently reported cybersecurity incidents, we also sought independent review and began a broader investigation of past model activity. We continue to [release our findings](#) as that investigation proceeds. We know that public confidence depends on the effectiveness of our protections, our openness about shortcomings, and the actions we take when something goes wrong.

4. Working with the government to strengthen safety and accountability.

Entering this new chapter in AI capabilities, we know that no company, industry, or government can meet this challenge alone. Democratic governments must ultimately determine the rules, safeguards, and accountability mechanisms.

Enforceable standards, independent scrutiny, and public accountability are necessary to ensure that increasingly capable AI is developed and used responsibly. This is why we have called for stronger rules on frontier AI safety. We support UK legislation to establish durable, mandatory, capability-based requirements for frontier AI, building on the existing voluntary cooperation between frontier labs and government. These requirements should focus on the small number of developers building the most advanced systems, with obligations proportionate to advanced systems' capabilities and risks. The priority should be serious national-security and frontier risks, including advanced cyber capabilities, biological risks and loss of control.

We have also called for [international cooperation](#) on shared standards to guide development across labs and countries. AI development and deployment are happening in multiple countries, adoption is occurring globally, and its impacts will likely be experienced by everyone in some way on some timeline. Democratic nations should therefore work together to develop compatible safety frameworks, trusted channels for information sharing, and coordinated responses to serious incidents.

OpenAI has also been an inaugural partner of UK AISI since its inception in 2023, sharing models with UK AISI for pre-deployment testing, and drawing on its expertise to [refine our systems](#), including our latest Astra model. We support the work of the International Network for Advanced AI Measurement,

OpenAI

Evaluation and Science and the UK's leadership of the Network, which is driving critical work on [international AI safety standards](#). We welcome the Prime Minister's ambition to pursue international collaboration on AI as part of the UK's G20 Presidency.

From Alan Turing to Bletchley Park and today's AI ecosystem, the UK has repeatedly helped shape transformative technologies. It can build on that heritage to lead again, leveraging its long-standing leadership in AI safety to strengthen public trust, drive international cooperation, and help ensure increasingly capable AI benefits society.

We share this Committee's commitment to ensuring AI expands economic opportunities for people in the UK. We look forward to discussing these issues in more detail at the upcoming evidence session, and to working closely with the UK Parliament to shape the future of AI.