

Opinion / World Opinion

## Opinion | Would China agree to slow AI development? That's the wrong question

Meaningful US-China action on frontier AI requires two things: a safe space for technical experts to collaborate and pressure on leading companies to act

**Xiao Qian**

Published: 8:30pm, 6 Oct 2026 / Updated: 8:30pm, 6 Oct 2026




Illustration: Mario Cameira

The latest [US-China summit](#) covered a wide range of topics including artificial intelligence (AI). However, the official readouts from both countries did not touch upon an issue that has generated heated debate: whether the United States and China should agree to pause or slow down the development of frontier AI.

That absence has triggered considerable discussion. During my recent visit to New York, many US friends asked whether Beijing would ever agree to some form of arrangement. Before predicting whether China would say yes or no, we need to unpack several other questions first.

The first question we need to clarify is who we are talking about when we ask, “What does China think?” If you ask people in China’s AI safety community, I expect many of them to be seriously concerned about frontier risks and already working on these questions. We can see this from the AI Safety Governance Framework 3.0 recently published by the National Technical Committee 260 on Cybersecurity of the Standardisation Administration of China.

If you ask Chinese AI companies struggling to survive and facing restrictions on access to advanced computing power, that question might seem rather distant from their immediate reality. If you ask ordinary Chinese people, many of whom still [see technological progress](#) as a way to improve their lives and worry more about being left behind, talk of existential risks and pausing frontier AI might sound like something from a sci-fi movie. There is no one firm answer from “China”.

The second question is more basic. What exactly counts as frontier AI?

According to recent estimates based on publicly released models and benchmark performance, the leading Chinese frontier models remain around eight months to one year behind those in the US. There is also concern in China that as US restrictions on access to advanced chips and computing power become stronger, this gap will become ever bigger.

---

 Item not supported in this print-friendly version

---

The capability gap changes the debate. While OpenAI's newly launched GPT-6 Astra has sparked renewed discussion over whether artificial general intelligence (AGI) has arrived, Chinese frontier labs still struggle to use limited compute to develop the best possible models while reserving enough resources for safety evaluations.

Z.ai recently delayed the public release of the weights of its GLM-5.3 model after finding that its cybersecurity capabilities required additional testing and safeguards. It shows that some Chinese developers are already confronting the trade-off between rapid release and safety evaluation.

The third question is more technical. What exactly would constitute a slowdown? Would limits apply to training runs, compute, model capabilities or deployment? How could compliance be verified? Who would have both the technical capacity and the political legitimacy to conduct such verification? Until these questions are answered, it is not realistic to expect a definite answer from any government to the larger question of whether it would agree to slow down.

There are also several aspects of the Chinese context that need to be understood. In China's modern history, technological development has usually been associated with economic growth and improvements in everyday life. For many Chinese people, the fear is more that they will be left behind and not benefit from advancements in technology.

In addition, Chinese discussions about technology have traditionally put the emphasis on balance. There is an old Chinese saying, *guo you bu ji*, which basically means that going too far is no better than falling short. The dominant tone of policy language is about balancing AI development and security, innovation and governance. The idea that technological development itself should be stopped because of a future catastrophic scenario seems to be a philosophical dilemma that has much less resonance in mainstream Chinese public debate.

It is also true that the discussion cannot be separated from the wider deterioration in US-China technological relations. Years of US restrictions on advanced chips, semiconductor equipment and sanctions on Chinese companies have created a perception in China that Washington is seeking to constrain China's technological development.

---

 Item not supported in this print-friendly version

---

Without fully understanding the complexity of US domestic politics and the differences among the Trump administration, Congress, AI companies and other stakeholders, Chinese observers can easily ask if “pacing” is another instrument for maintaining US technological superiority while slowing down China.

Rather than waiting for a definite answer from China, perhaps we should think about what can be done to create an environment in which governments are more ready to take action when the question becomes more urgent.

As China and the US prepare for the next round of intergovernmental dialogue on AI, frontier risks can be part of the agenda. There should also be a safe space for technical experts from both countries to work on possible governance frameworks. During the Cold War, the Pugwash Conferences on Science and World Affairs created space for scientists from different sides of the geopolitical divide to work through some of the most difficult questions surrounding nuclear weapons and arms control.

AI is obviously different from nuclear weapons, but the underlying idea remains useful. When official negotiations are difficult, scientists and technical experts need places where they can explore what is technically possible before governments are ready to negotiate.

What is more important is that frontier companies need to be encouraged to begin by taking measures themselves, including investing more in safety evaluations, developing clearer release thresholds, testing models for loss-of-control risks and publishing evidence about which safeguards actually work. If Chinese and American companies can compete over who develops better safety evaluations, stronger safeguards and more credible solutions, that will be the kind of competition that people welcome.