AI Slowdown Is Not Necessarily a Negative
Understanding the Logic Behind the Latest Signals from AI Leaders

Over the weekend, Dario Amodei, CEO of Anthropic, called for a global slowdown in the pace of frontier AI development to allow more time for AI safety research. Elon Musk expressed his support by reposting the statement, while OpenAI CEO Sam Altman noted that moderating the pace of frontier AI development has recently become an important topic of internal discussion at OpenAI. Meanwhile, reports that OpenAI may postpone its IPO until next year have raised concerns among investors about the outlook for the AI industry and the valuation of AI-related assets.

We believe, however, that these developments do not necessarily represent a negative for the AI supply chain as a whole. Rather, they may signal a shift in the focus of AI development.

First, the core of the current discussion is safety governance during the training of frontier models, rather than restrictions on the actual use of AI.

Real-world usage data continues to point to rapid growth in inference demand. Following the launch of OpenAI's next-generation model, Astra, data tracked by OpenRouter showed that OpenAI's weekly Token usage increased by approximately 31%. OpenAI also temporarily suspended Pro subscriptions, not because of weak demand, but because user demand exceeded expectations and available computing capacity was unable to keep pace with usage.

The current industry dynamic appears to be that once large language models reach a sufficient level of usability, real-world adoption can expand rapidly. A slower pace of new model releases does not necessarily translate into weaker demand for AI services. What leading AI companies are seeking to moderate is primarily the pace of the model capability race, while underlying demand from end users and commercial applications continues to increase.

Second, based on the industry developments we are currently observing, leading model developers have not stopped training. Instead, they are moving alignment-related work earlier in the development process, with additional alignment and verification procedures being introduced before, during, and after training.

The pace of public model releases may slow to some extent going forward. More human intervention may also be incorporated into the training process, while some companies may deploy two separate models to conduct mutual verification and review.

In other words, beyond the computing resources required for model training itself, additional computing capacity may be required for review and validation. This could actually create incremental demand for computing resources.

Such changes in the development process do not necessarily reduce overall training-side compute consumption. Training, supervised fine-tuning (SFT), alignment and debugging, as well as retraining following interruptions, all continue to consume substantial computing resources.

At present, computing demand at model developers such as OpenAI and Anthropic is roughly split between training and inference, with each accounting for around half of total consumption. For large cloud service providers, however, inference already represents a higher share of computing demand. This indicates that inference is gradually becoming a larger component of overall AI compute demand.

Moreover, when computing resources are relatively constrained, a temporary moderation in training demand could allow more capacity to be redirected toward inference.

Third, historical experience suggests that when major technology companies call for a slowdown, this does not necessarily mean that the broader industry cycle has peaked.

For example, in April this year, Anthropic released its Mythos model and established the Project Glasswing alliance, continuing to promote the development of a more comprehensive AI safety framework across the industry. At the same time, assets across the global AI supply chain continued to perform strongly, with semiconductors and computing infrastructure maintaining relatively high levels of industry momentum.

This suggests that the market is placing greater emphasis on the long-term growth potential of the AI industry rather than short-term changes in the pace of model releases.

Finally, regarding recent discussions surrounding the postponement of OpenAI's IPO, we believe this is more likely to reflect the timing of the listing process than a change in the underlying industry trend.

Objectively, completing an IPO within 2026 would already face a number of practical constraints. Anthropic is reportedly planning to begin its IPO roadshow in October, the U.S. midterm elections will take place in November, and December marks the year-end holiday period in both Europe and the United States. This leaves OpenAI with a relatively limited window for completing an IPO this year.

Against this backdrop, a postponement of the IPO may be more reflective of practical scheduling considerations and should not be overinterpreted as a signal of a fundamental change in the AI industry outlook.

Overall, we believe that the current discussion around AI safety governance is unlikely to alter the long-term direction of AI development. Instead, it may accelerate the industry's transition from a phase focused primarily on model capabilities toward a new stage in which capabilities, safety, and commercialization develop in parallel.

The growth drivers of the AI industry are also gradually shifting from infrastructure build-out toward rising inference demand and the expansion of the application ecosystem.

We therefore remain constructive on the AI sector over the medium to long term. If market sentiment leads to significant short-term volatility, this could provide more attractive opportunities for investors to reassess the sector. Within the AI value chain, we continue to focus on structural opportunities, including continued upgrades in AI hardware, the development of inference infrastructure, and the accelerating adoption of AI applications and AI-

enabled enterprise workflows.

Disclaimer